



US008988426B2

(12) **United States Patent**
Chen et al.

(10) **Patent No.:** **US 8,988,426 B2**
(45) **Date of Patent:** **Mar. 24, 2015**

(54) **METHODS AND APPARATUS FOR RENDERING LABELS BASED ON OCCLUSION TESTING FOR LABEL VISIBILITY**

(75) Inventors: **Billy P. Chen**, Santa Clara, CA (US);
Jeffrey P. Hultguist, Cupertino, CA (US); **Christopher Blumenberg**, San Francisco, CA (US); **Aroon Pahwa**, Palo Alto, CA (US)

(73) Assignee: **Apple Inc.**, Cupertino, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 173 days.

(21) Appl. No.: **13/601,934**

(22) Filed: **Aug. 31, 2012**

(65) **Prior Publication Data**

US 2013/0321397 A1 Dec. 5, 2013

Related U.S. Application Data

(60) Provisional application No. 61/655,805, filed on Jun. 5, 2012.

(51) **Int. Cl.**
G06T 15/00 (2011.01)
G01C 21/36 (2006.01)
G06T 17/05 (2011.01)
G06T 19/00 (2011.01)

(52) **U.S. Cl.**
CPC **G06T 17/05** (2013.01); **G01C 21/3673** (2013.01); **G06T 19/00** (2013.01); **G06T 2219/004** (2013.01)

USPC **345/419**

(58) **Field of Classification Search**

None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,557,736 B1 7/2009 Daily et al.
7,609,258 B2 10/2009 Suzuno et al.

(Continued)

FOREIGN PATENT DOCUMENTS

WO 2011/095227 8/2011

OTHER PUBLICATIONS

Over et al, Generating web-based 3D City Models from OpenStreetMap: The current situation in Germany, 2010, Computers, Environment and Urban Systems 34, 496-507.*

(Continued)

Primary Examiner — Kee M Tung

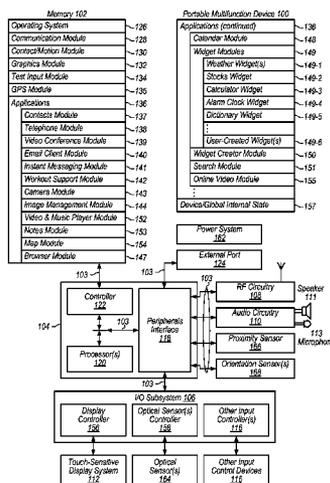
Assistant Examiner — Xin Sheng

(74) *Attorney, Agent, or Firm* — Robert C. Kowert; Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.

(57) **ABSTRACT**

Methods and apparatus for a map tool displaying a three-dimensional view of a map region, where the map tool determines whether or not partially occluded labels within the map region are to be drawn or are to not be drawn. The map tool determines whether or not to draw a label in a map view dependent upon mapping information and upon construction of a three-dimensional model based on one or more two- or three-dimensional data sets. The map tool further determines whether or not to draw a label in the map view dependent upon a measure of occlusion of a label in the map view. In order to determine a measure of occlusion, the map tool may calculate whether a line of sight projection from virtual camera viewpoint for the mobile device intersects, once or more than once, with any objects or landmarks within the three-dimensional model.

25 Claims, 17 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2010/0141648	A1 *	6/2010	Bell et al.	345/419
2011/0001751	A1	1/2011	Carlsson et al.	
2012/0019513	A1	1/2012	Fong et al.	
2012/0245848	A1 *	9/2012	Spindler et al.	701/532

OTHER PUBLICATIONS

Vittayakorn et al, Quality Assessment for Crowdsourced Object Annotations, 2011, BMVC.*
Peterson et al, Managing Visual Clutter: A Generalized Technique for Label Segregation using Stereoscopic Disparity, 2008, IEEE.*

* cited by examiner

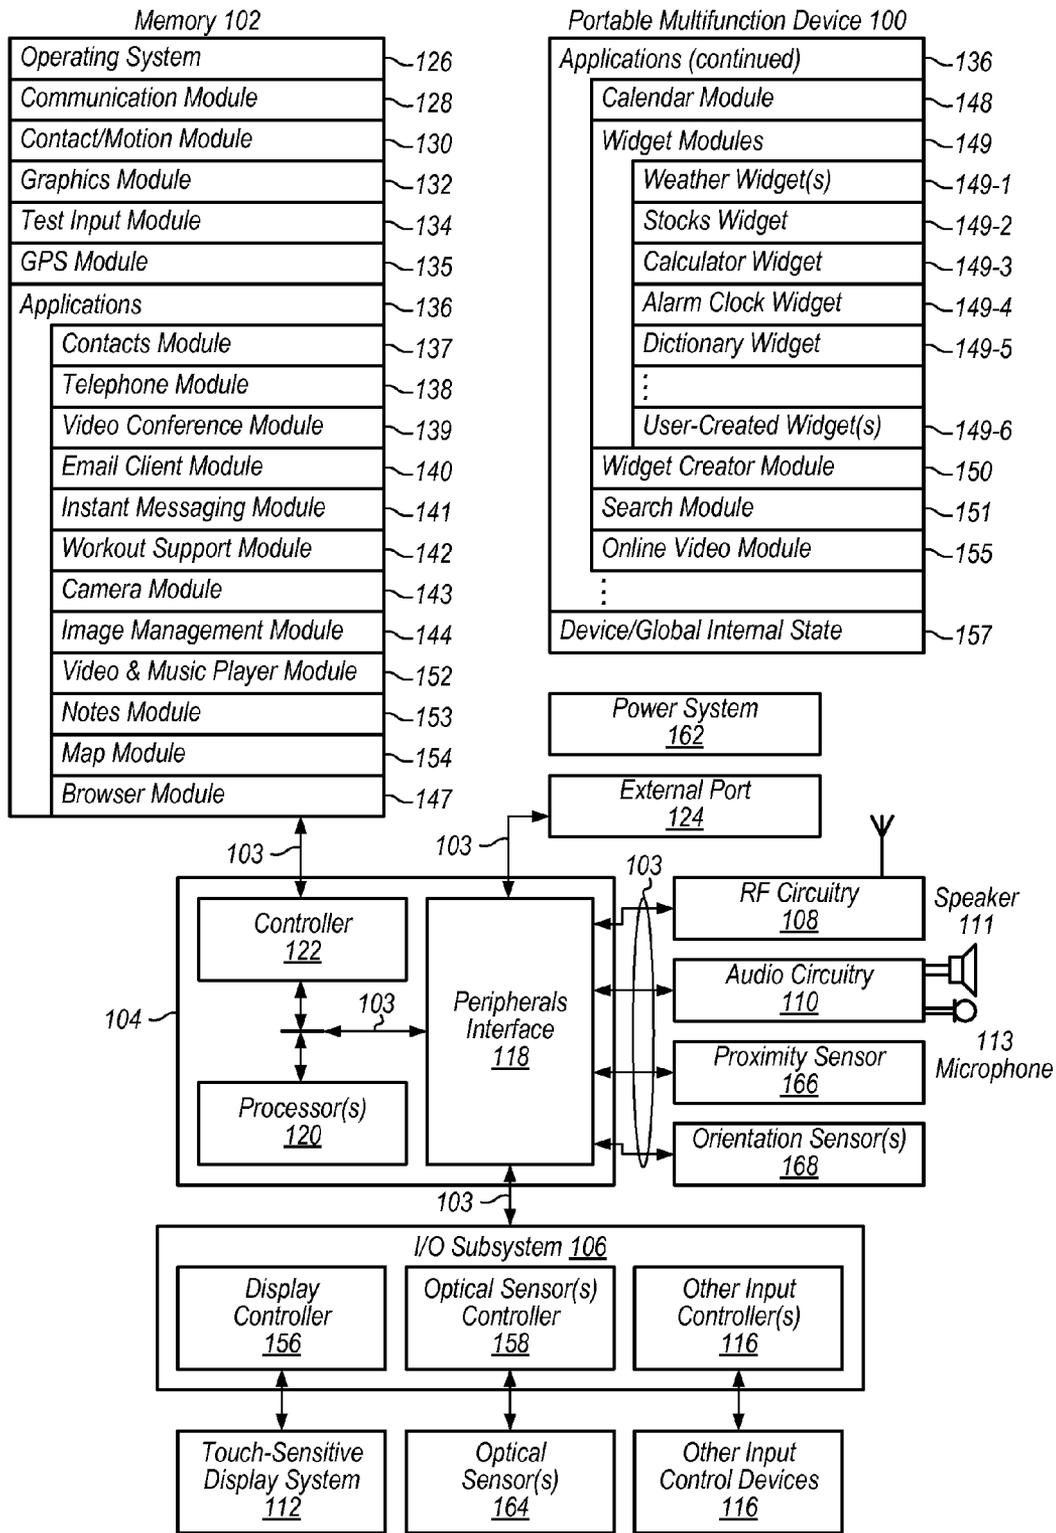


FIG. 1A

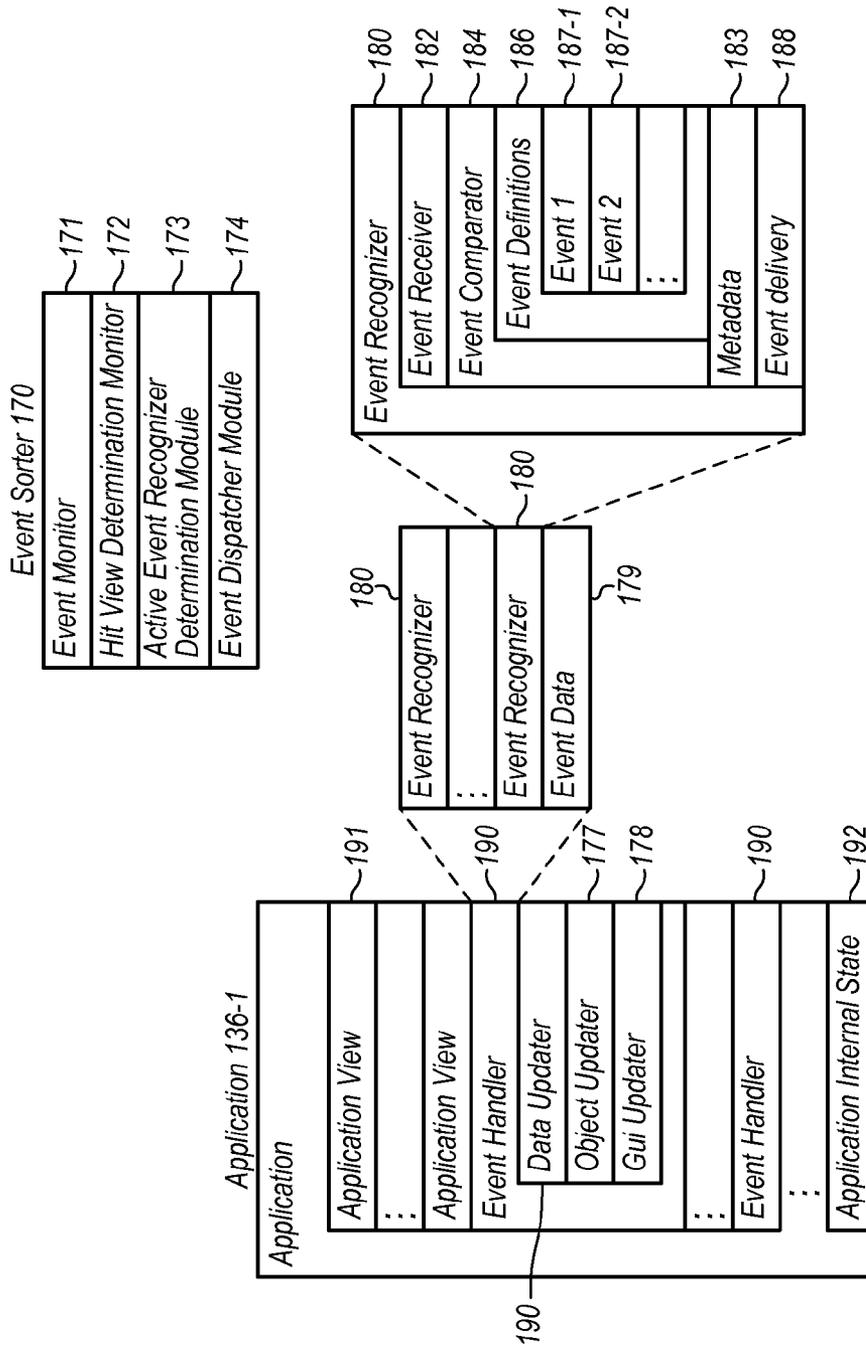


FIG. 1B

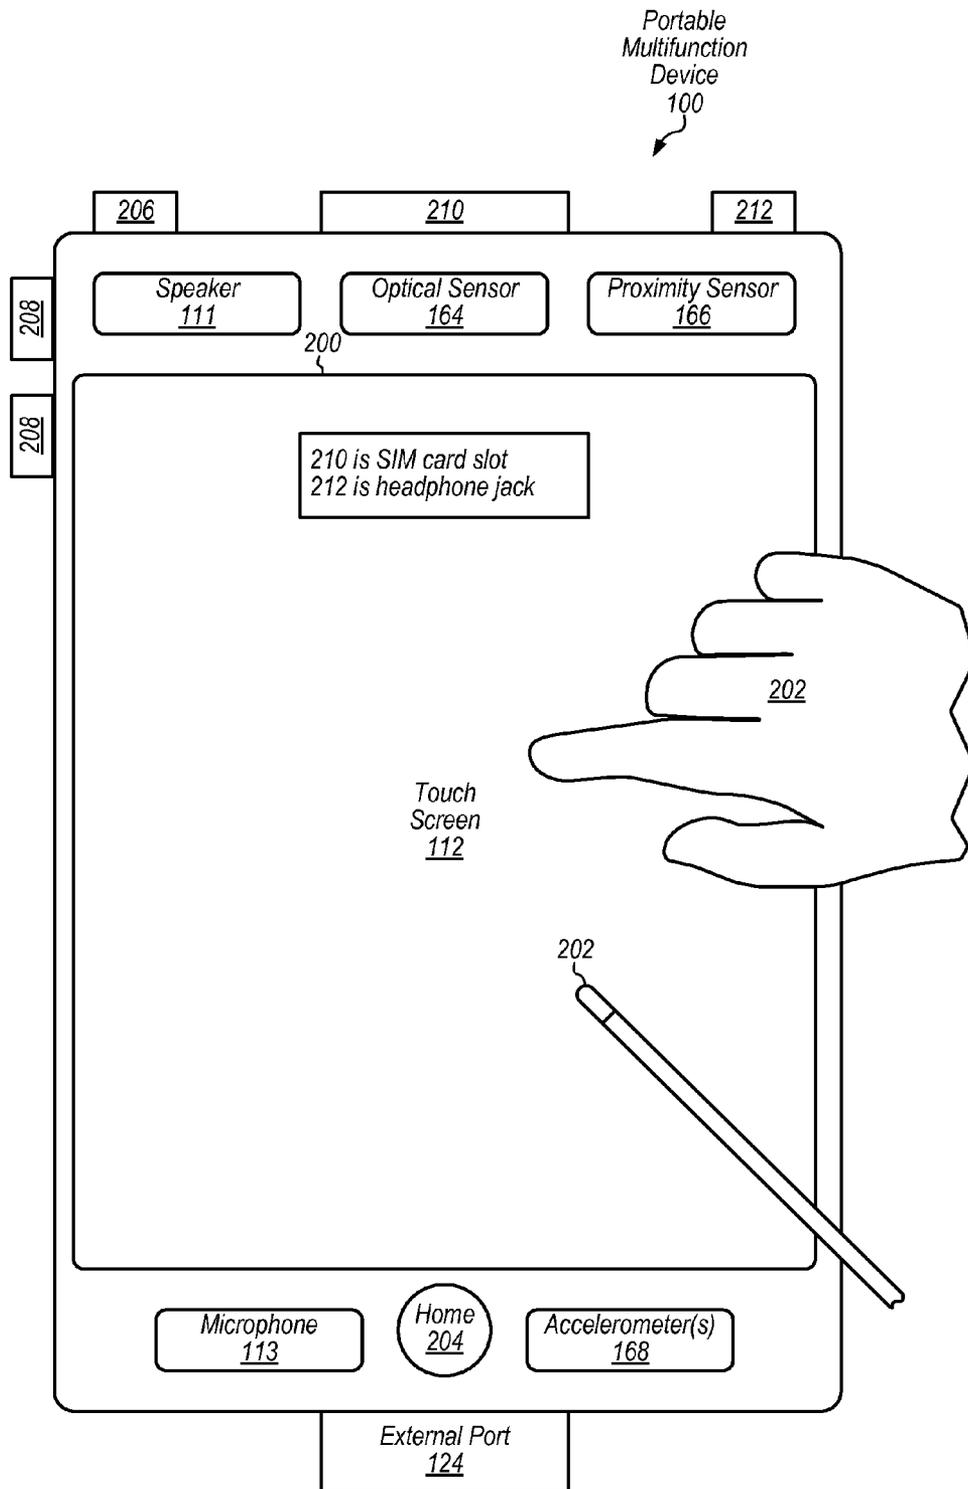


FIG. 2

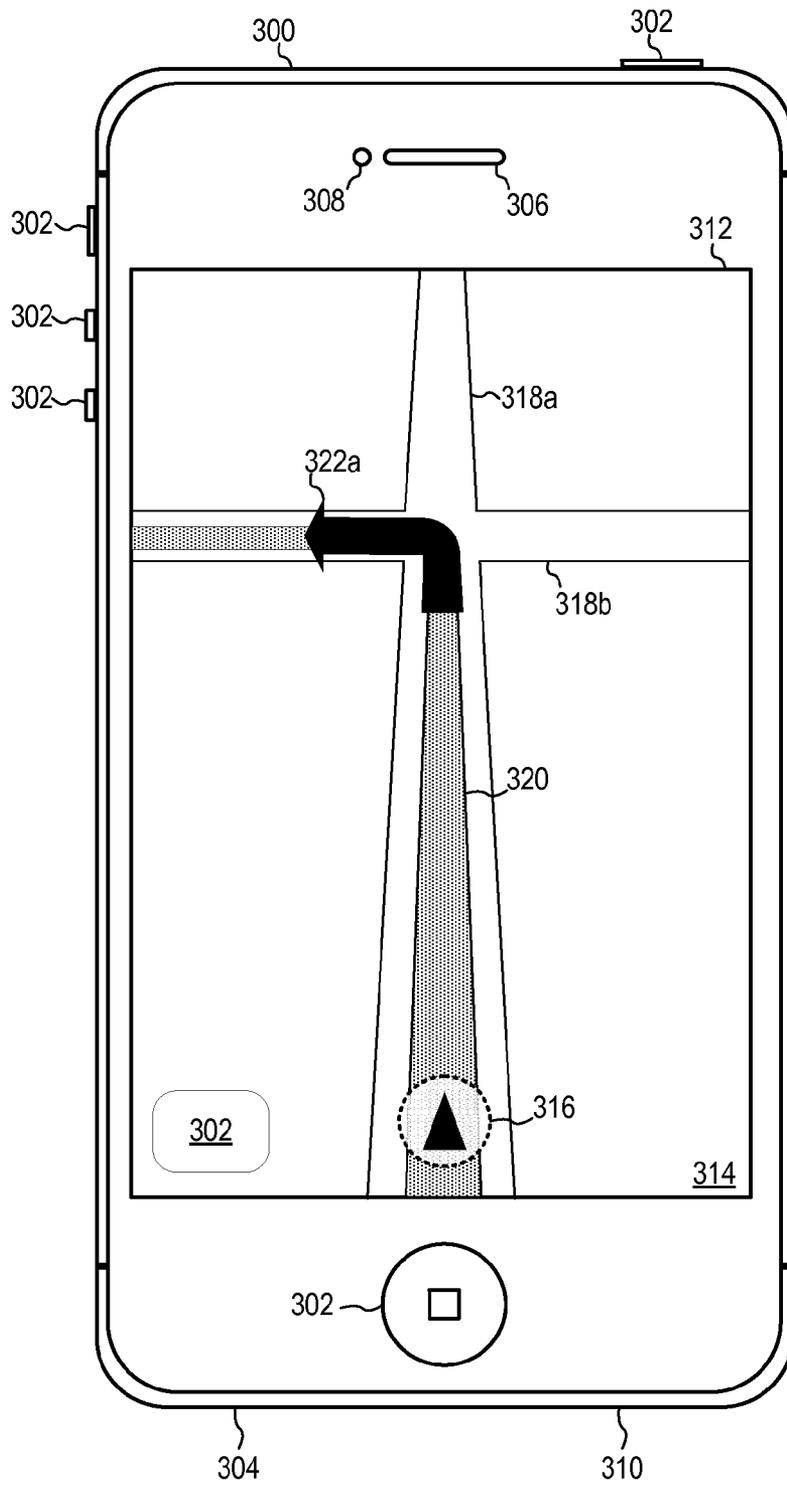


FIG. 3

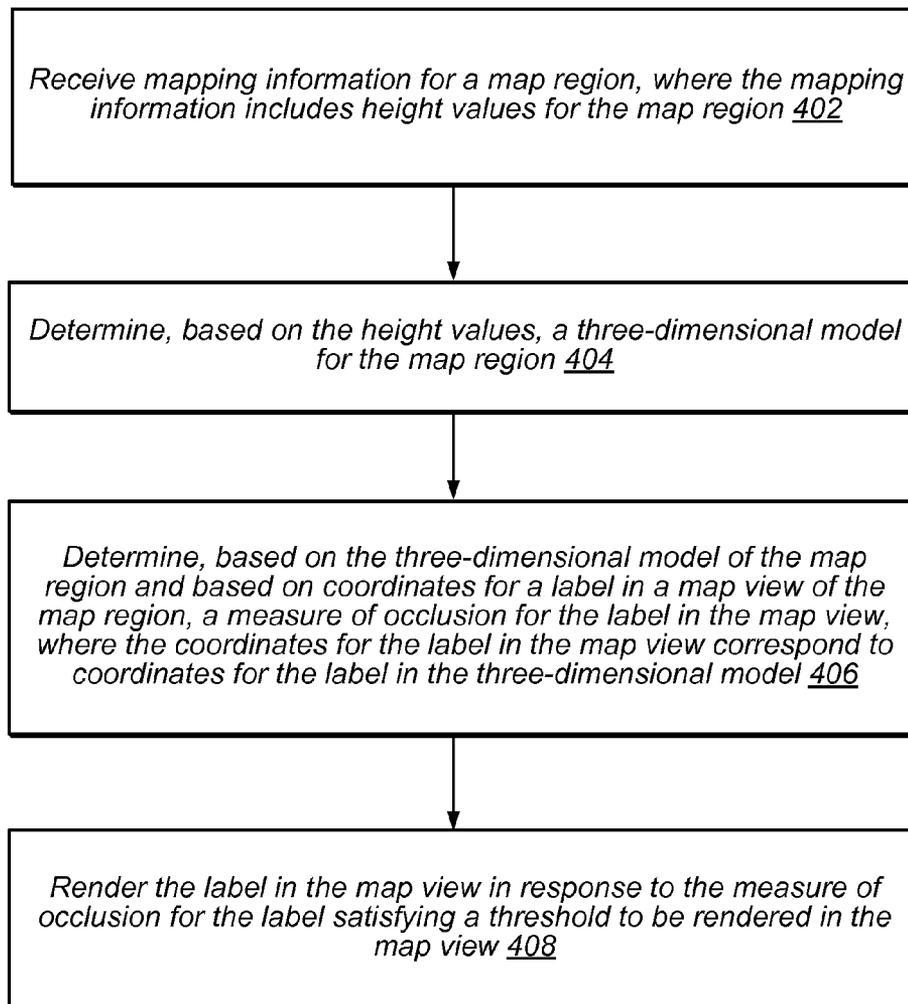


FIG. 4A

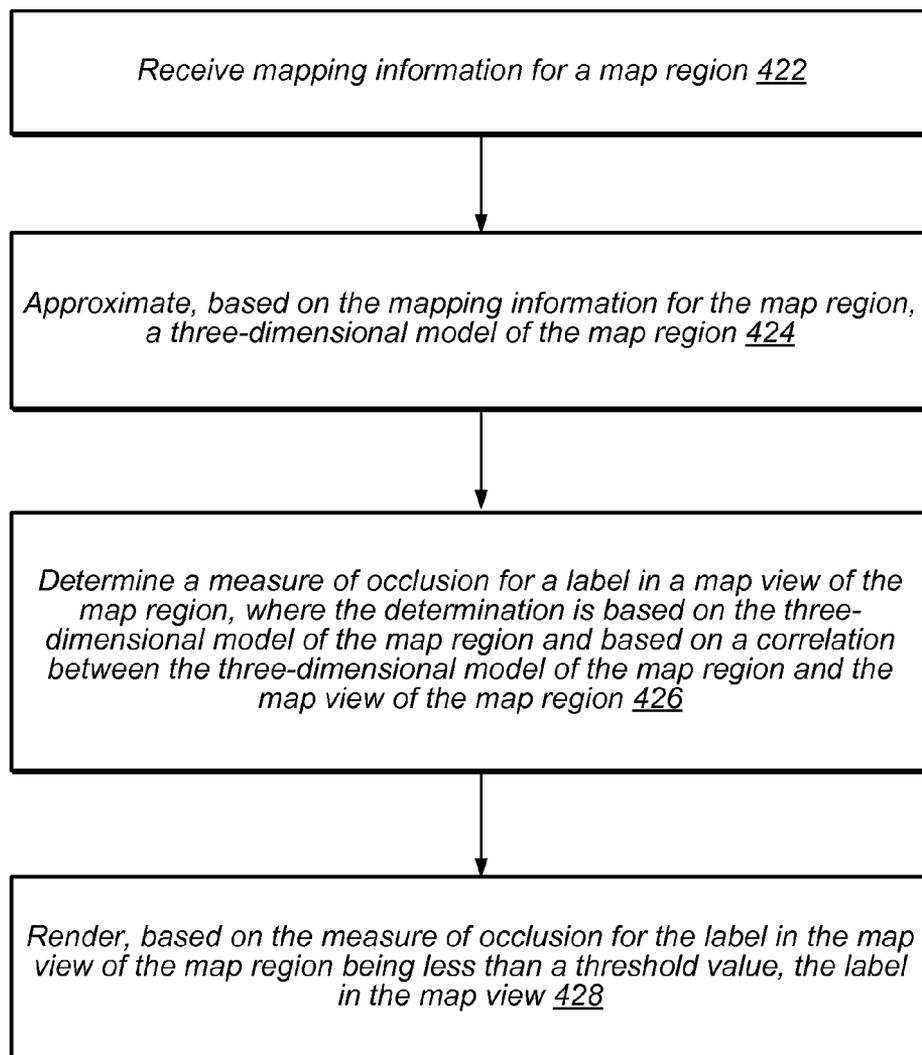


FIG. 4B

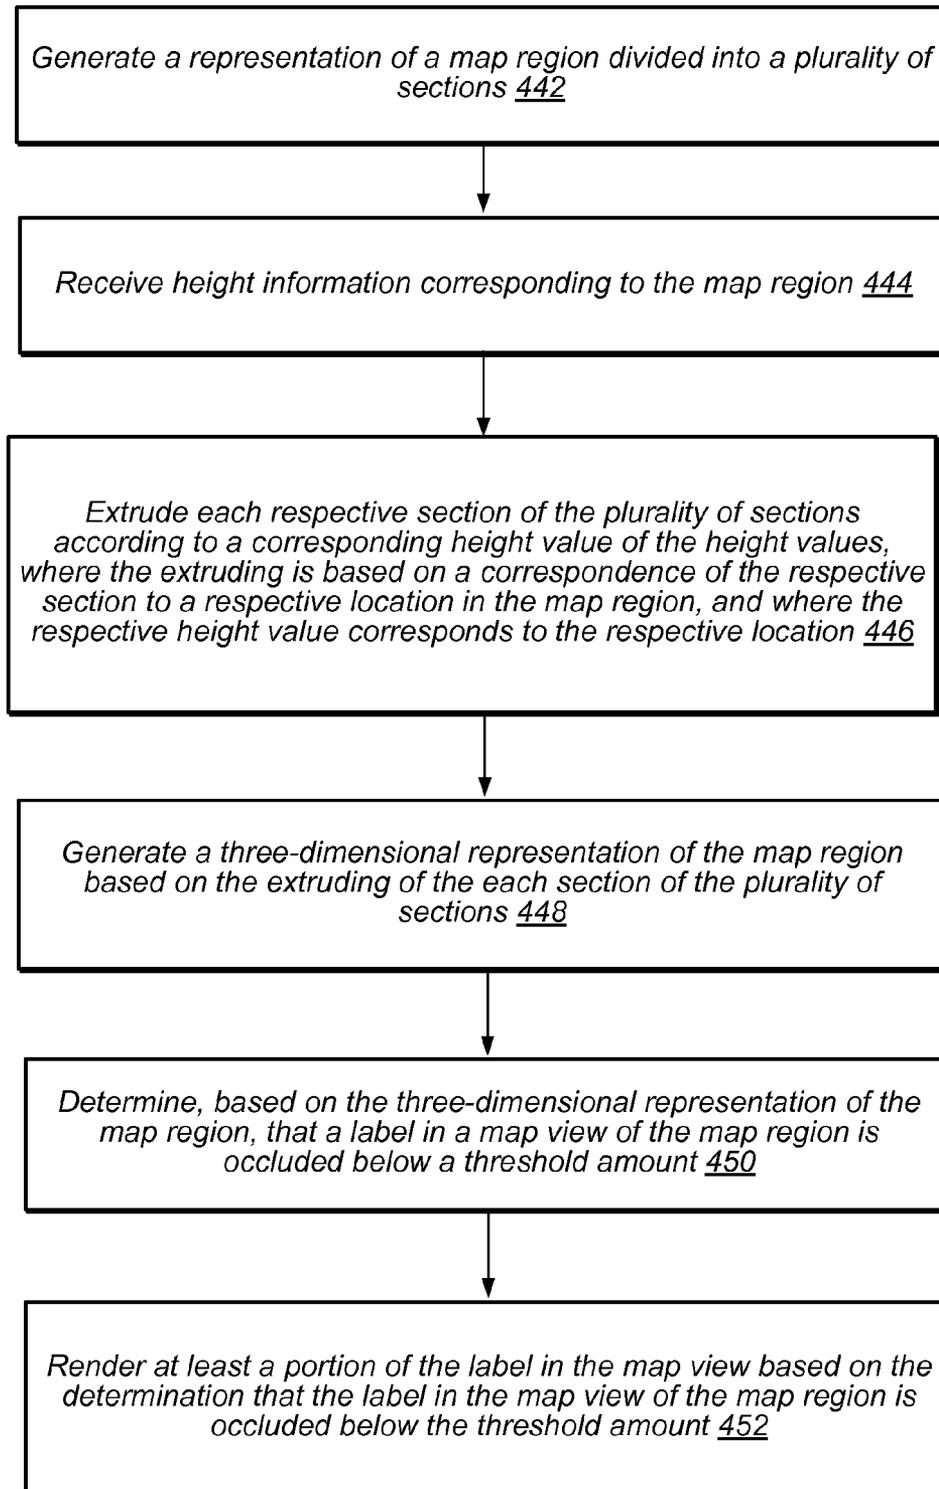


FIG. 4C

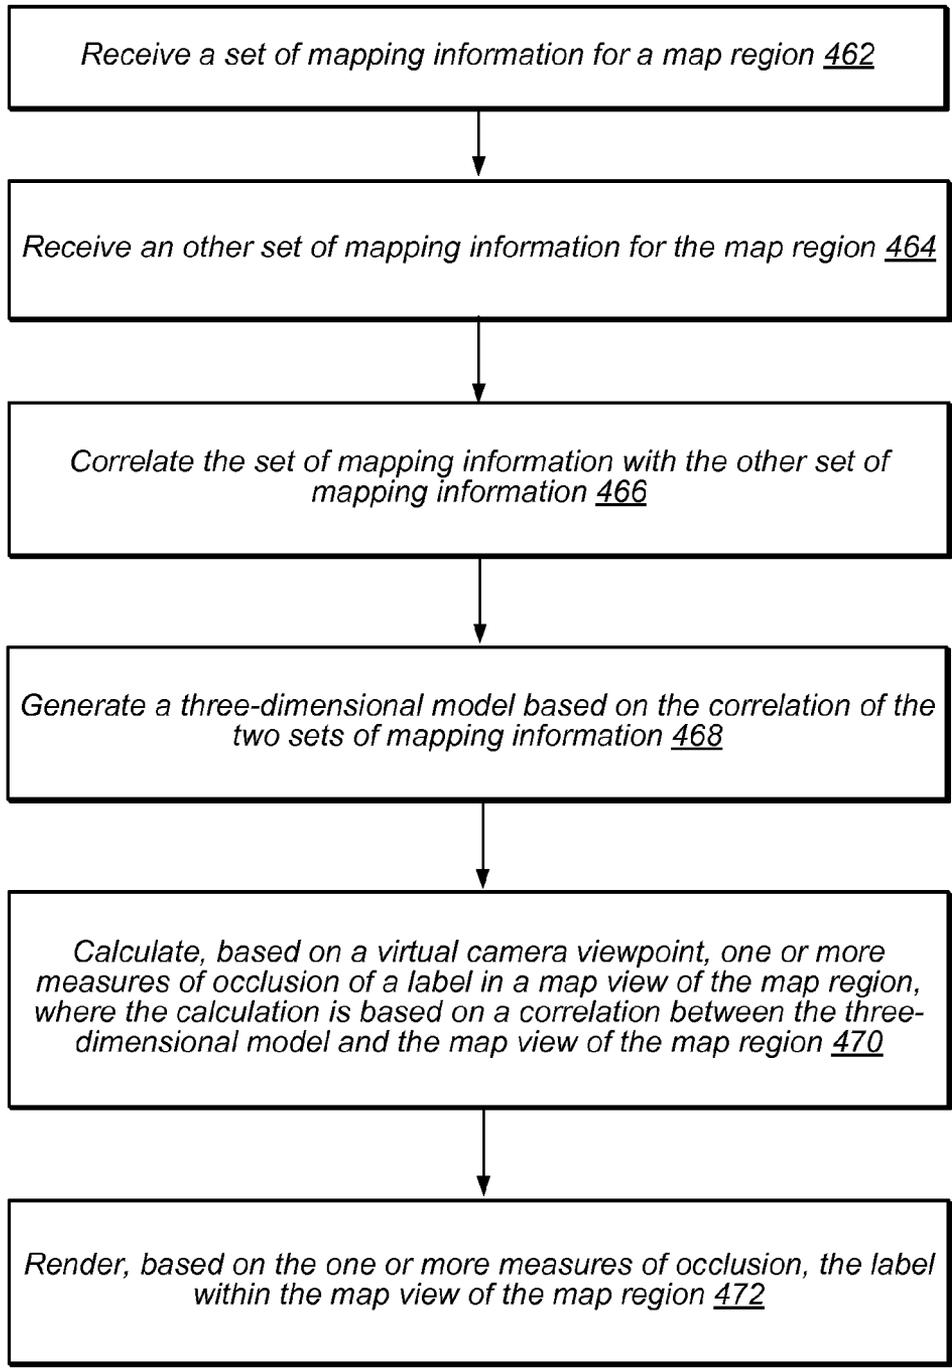


FIG. 4D

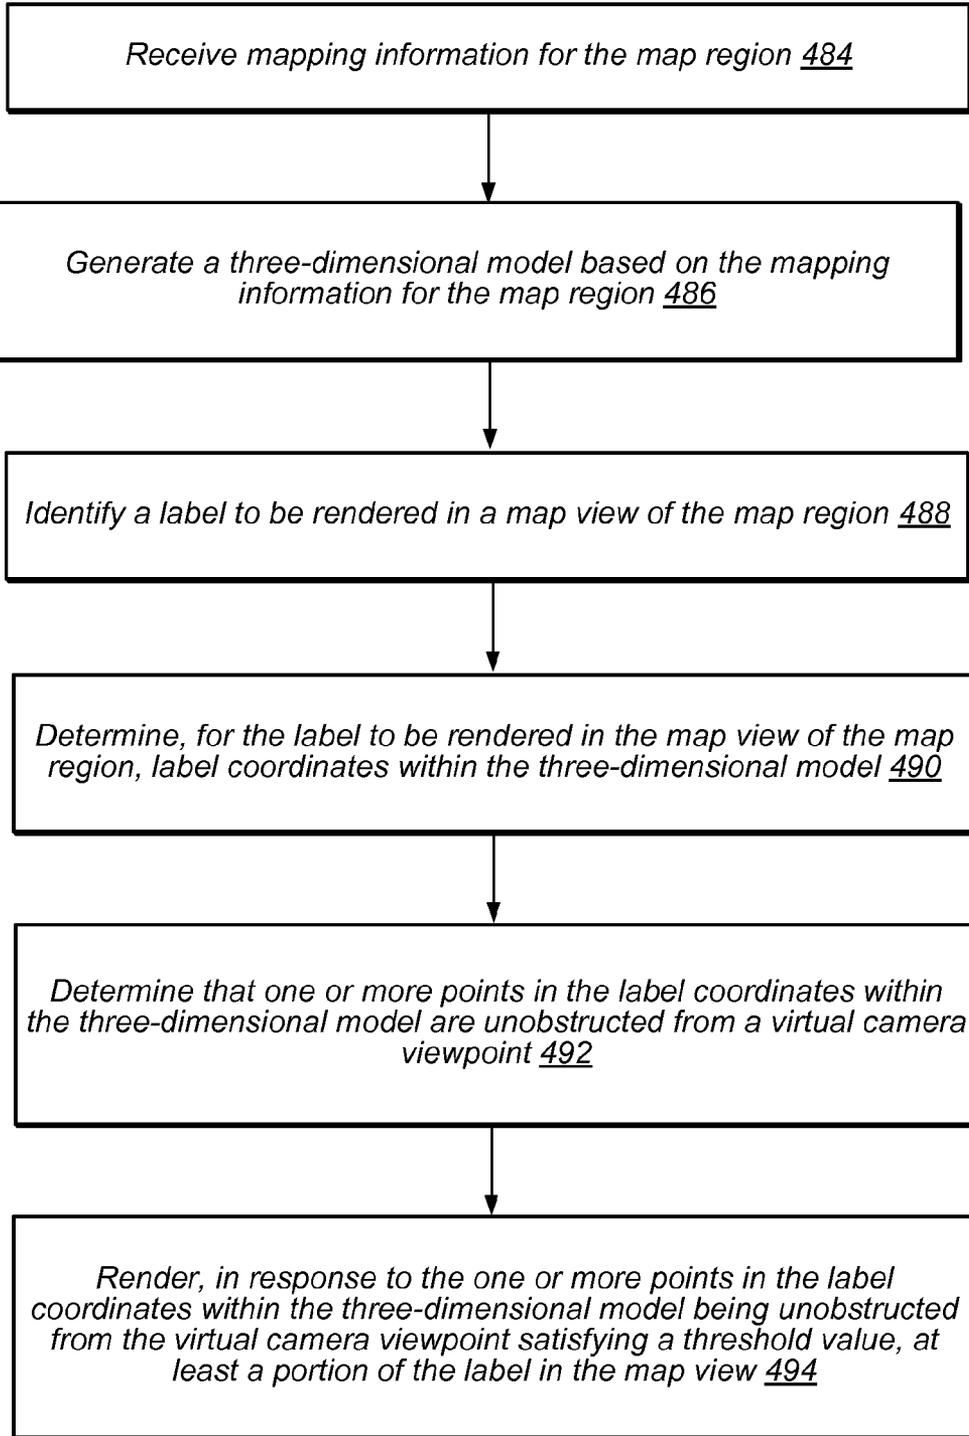


FIG. 4E

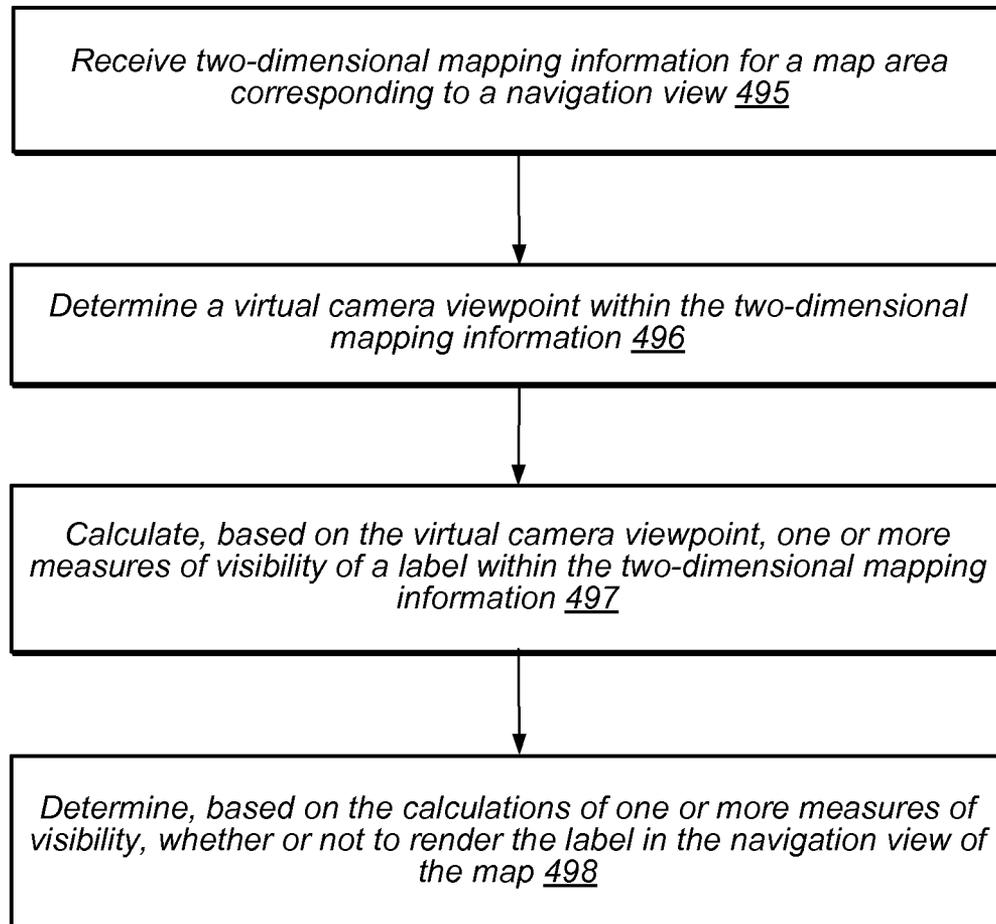


FIG. 4F

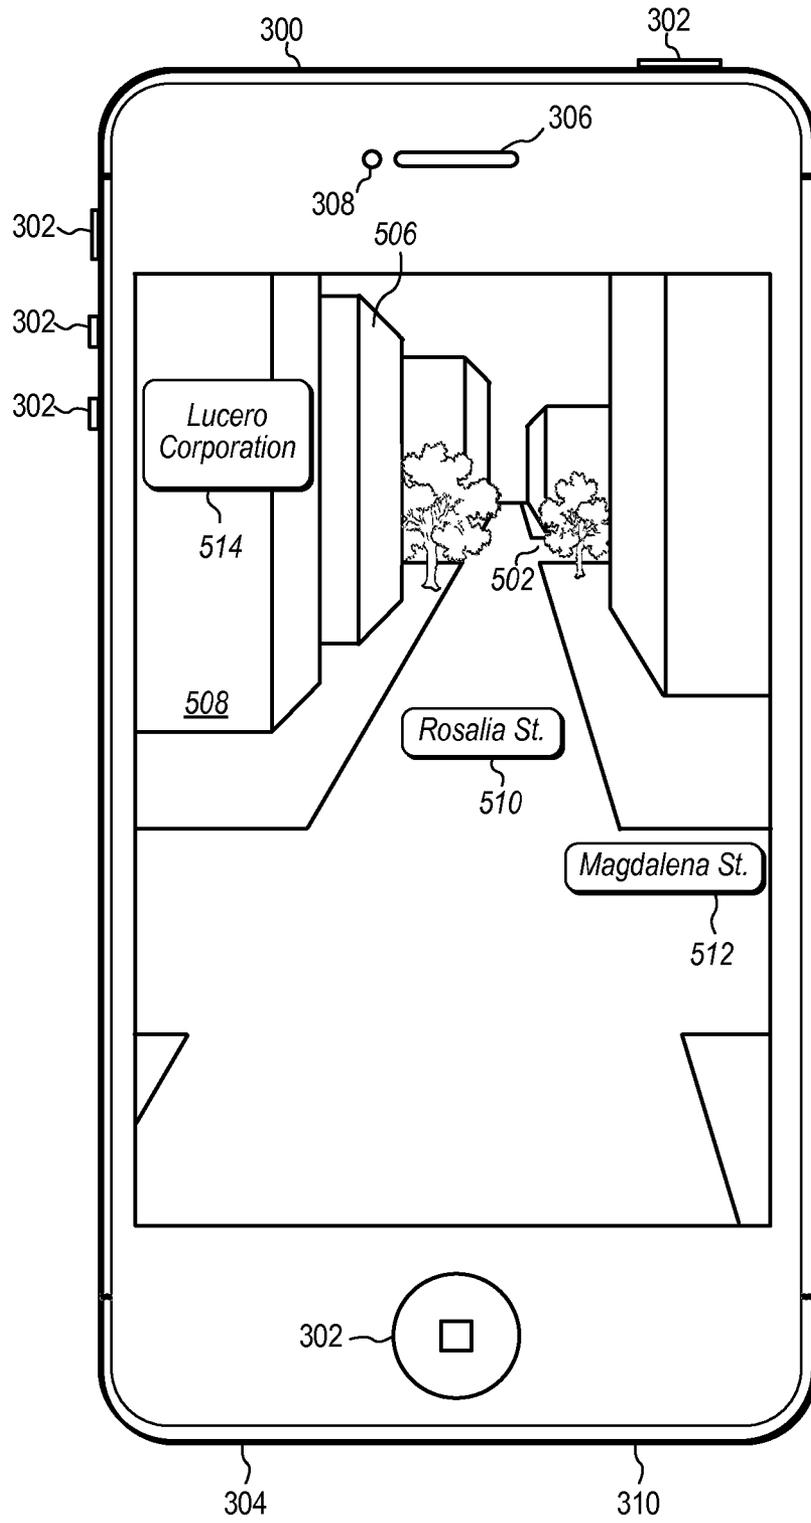


FIG. 5A

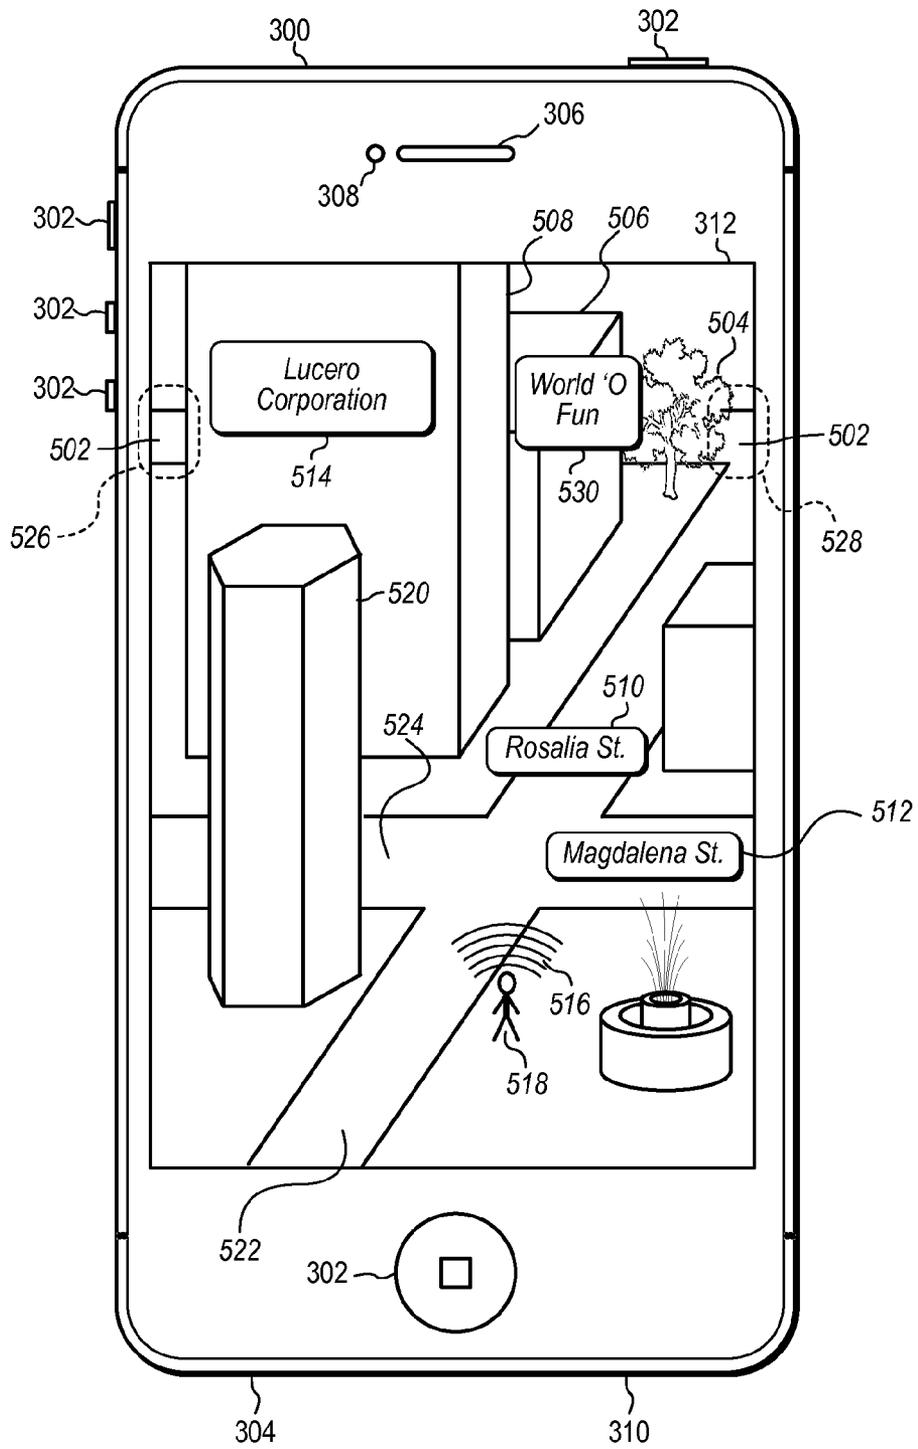


FIG. 5B

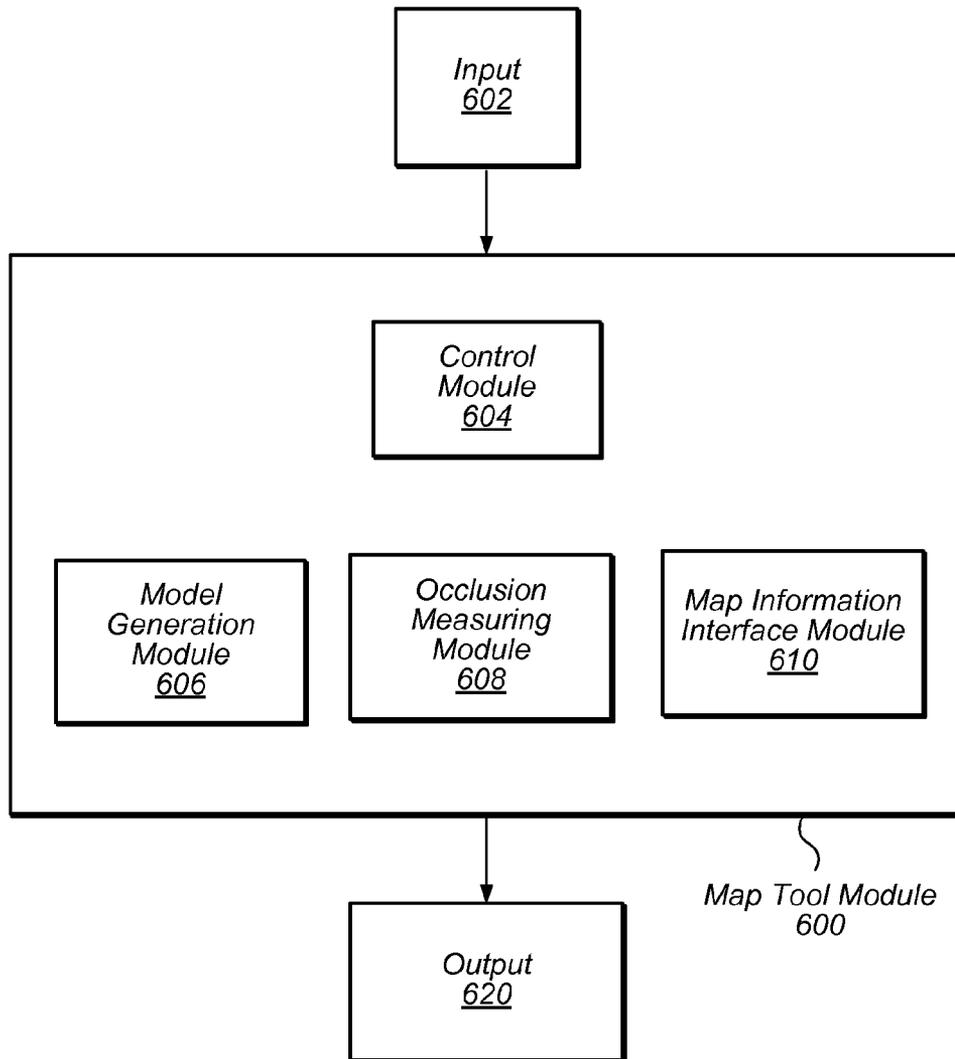


FIG. 6

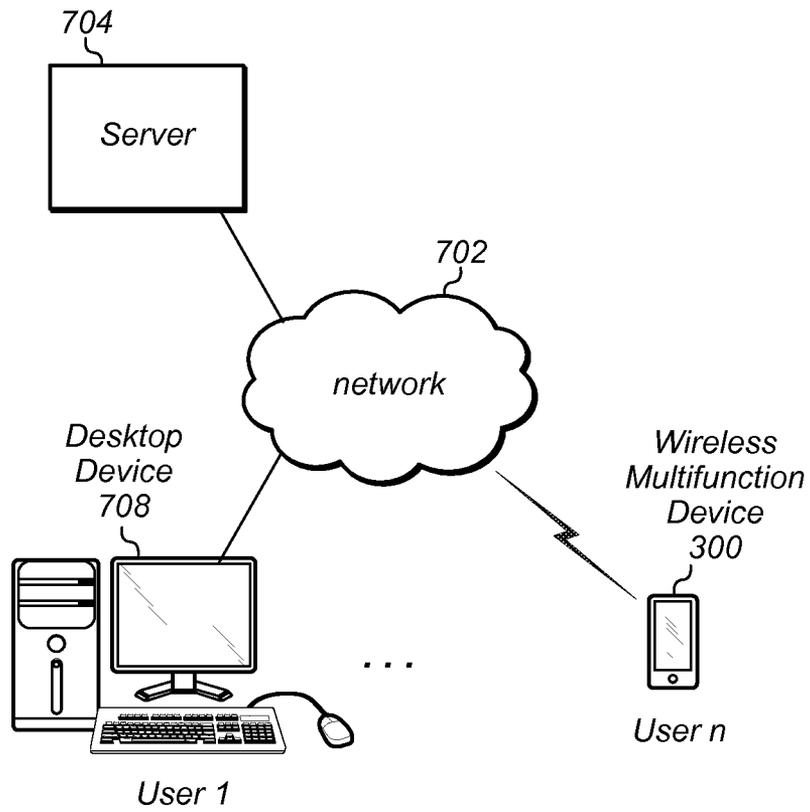


FIG. 7

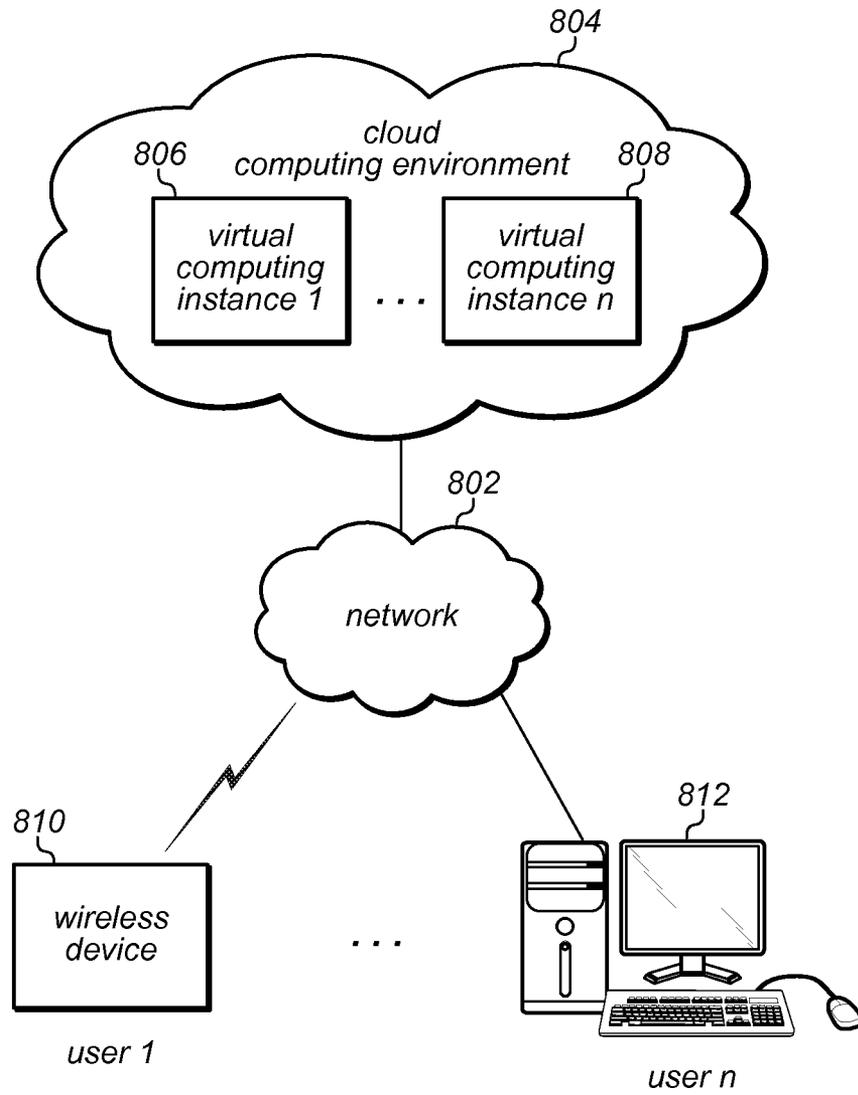


FIG. 8

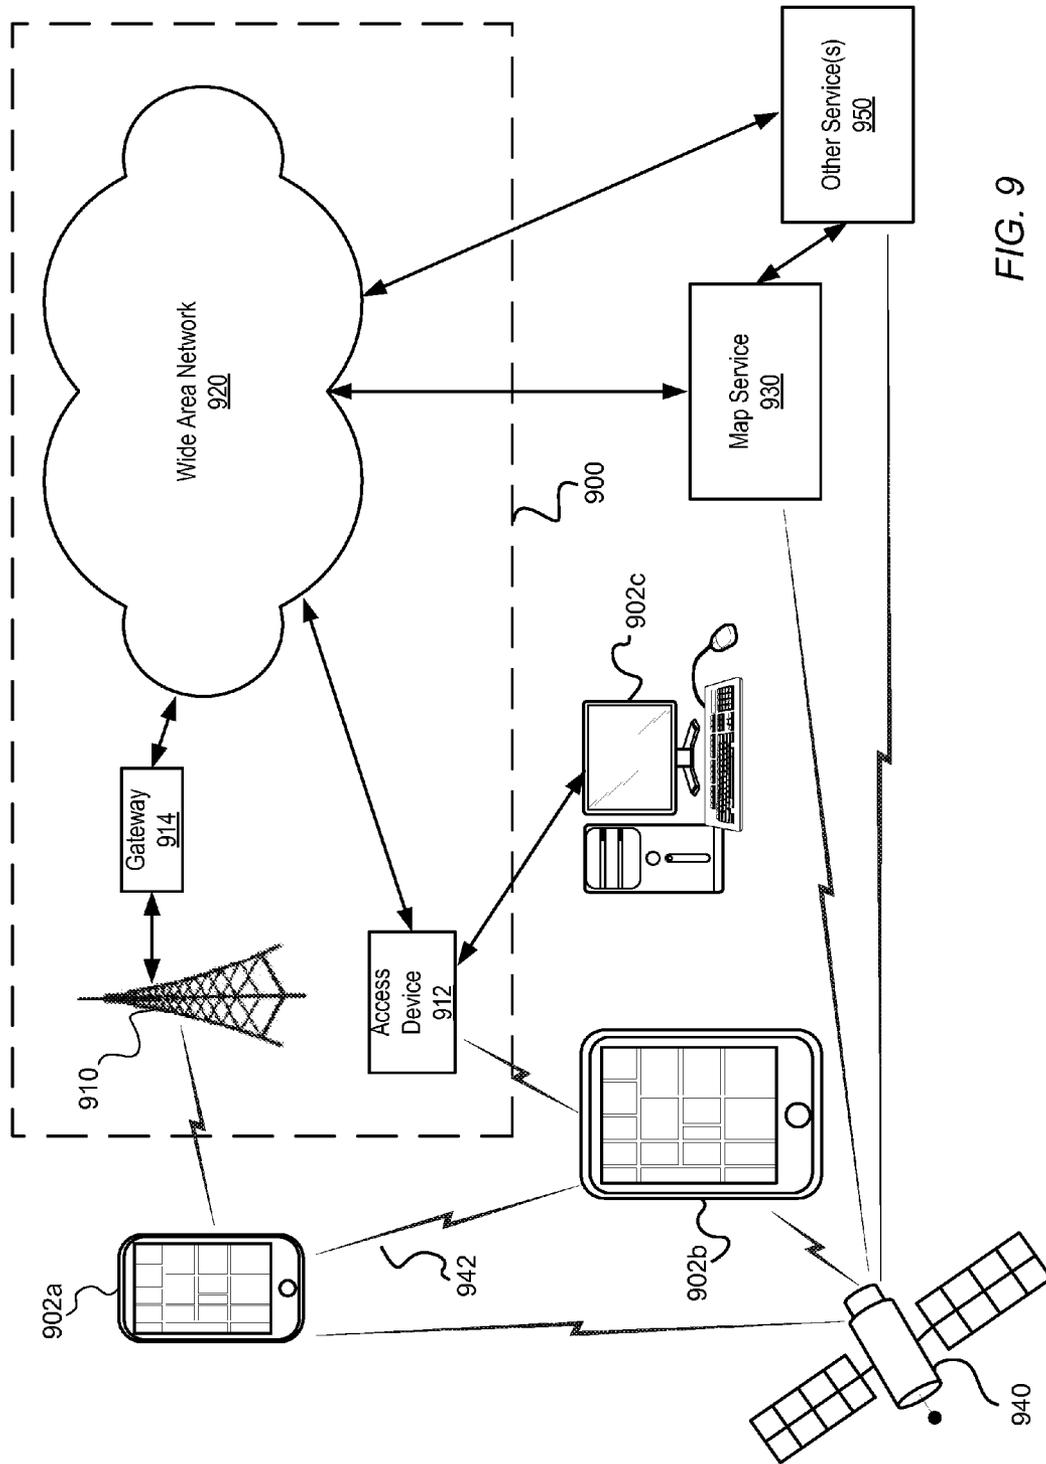


FIG. 9

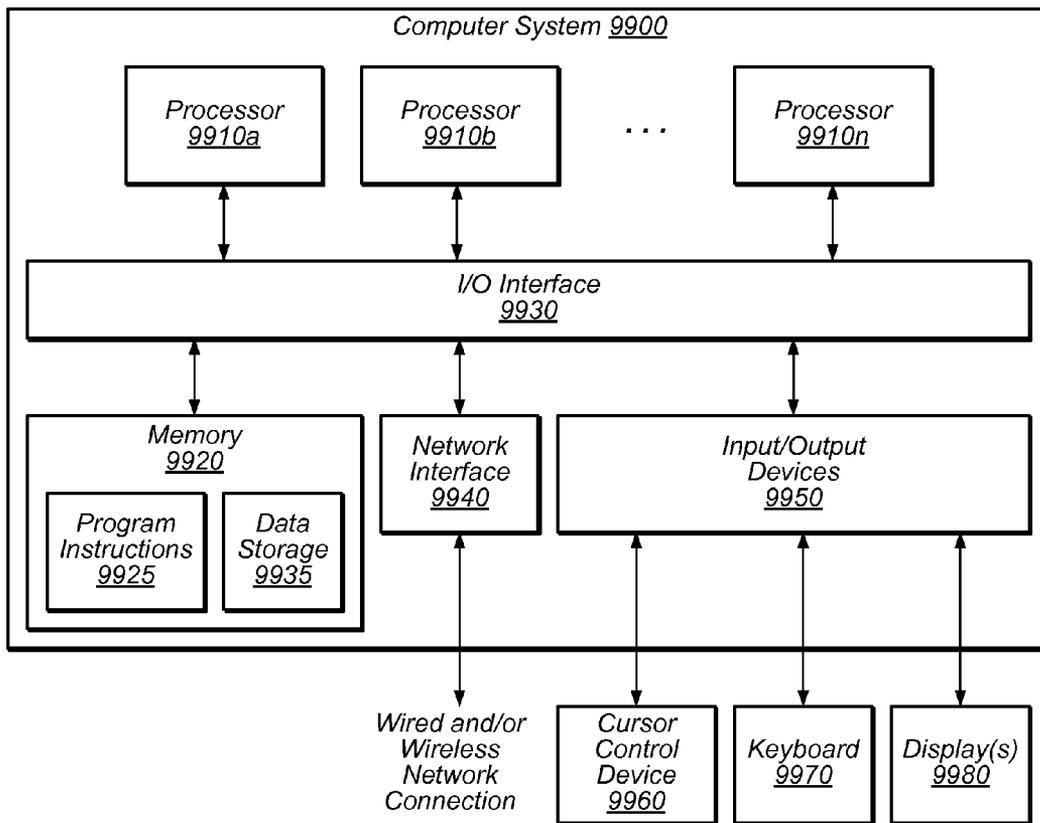


FIG. 10

METHODS AND APPARATUS FOR RENDERING LABELS BASED ON OCCLUSION TESTING FOR LABEL VISIBILITY

This application claims benefit of priority to U.S. Provisional Application Ser. No. 61/655,805, entitled "Methods and Apparatus for Rendering Labels Based on Occlusion Testing for Label Visibility," filed Jun. 5, 2012.

BACKGROUND

Mobile devices and desktop computers often provide various map related functions, including map views and navigation capabilities. Within a map view of a map application executing on a mobile device, one or more labels may be displayed corresponding to various landmarks or objects within the map view. For example, as a user walks down a street referencing the map view provided by a mobile device, street names may be displayed within the map view. In some cases, the map view may be three-dimensional, and in such a case some of the labels drawn within the map view may be partially or completely invisible due to obstructions such as buildings. In the case that a label is blocked, or occluded, to such an extent that the label provides no meaningful information, it may be preferable for the map application to draw the map view without the label at all. A positive effect from not drawing labels that provide little value is a cleaner and less cluttered version of the map view.

SUMMARY

In one embodiment, a mobile device may receive mapping information, which the mobile device may use to construct a three-dimensional model corresponding to a map region of a map view to be rendered on a display screen for the mobile device. A map tool on the mobile device may approximate one or more spatial dimensions for one or more of the objects in the map area to be displayed in the map view to determine whether or not to display a given label in the map based on whether the label is occluded or how occluded the label be if rendered.

For example, in the case a user is standing on a city street, the map tool may approximate the spatial dimensions of each building in the map view, where the approximations are based on mapping information. The map tool may then determine a measure of occlusion for a given label in the map view for the map, where the determination of the measure of occlusion is based on the previous approximations and on a virtual camera viewpoint.

In some cases, the measure of occlusion may be a function providing a result indicating whether enough of the label is visible in order for the label to be drawn in the map view. In some cases, the map tool may then determine whether the measure of occlusion is greater than an adjustable threshold value. In response to the measure of occlusion being greater than the threshold value, the map tool may then render the label in the map view.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is an illustration of a portable multifunction device suitable to implement a map tool, according to some embodiments.

FIG. 1B is a diagram illustrating example components within a portable multifunction device suitable for implementing a map tool, according to some embodiments.

FIG. 2 illustrates a touch screen on a multifunction device, according to some embodiments.

FIG. 3 illustrates another multifunction device configurable to implement a mapping application and map tool, according to some embodiments.

FIGS. 4A-4F depict example flowcharts corresponding to different embodiments of a map tool, according to some embodiments.

FIGS. 5A and 5B depict illustration different viewpoints available in a map view, according to some embodiments.

FIG. 6 illustrates map tool module, according to some embodiments.

FIG. 7 depicts an illustration of a client/server architecture that may be used in some embodiments of the map tool.

FIG. 8 depicts an illustration of a cloud computing environment that may be used in some embodiments of the map tool.

FIG. 9 depicts elements of a map service operating environment, according to some embodiments.

FIG. 10 depicts elements of an example computer system capable of implementing a map tool.

While the invention is described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the invention is not limited to the embodiments or drawings described. It should be understood that the drawings and detailed description are not intended to limit the invention to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention. The headings used are for organizational purposes only and are not meant to be used to limit the scope of the description. As used throughout this application, the word "may" is used in a permissive sense (meaning "having the potential to"), rather than the mandatory sense (meaning "must"). Similarly, the words "include", "including", and "includes" mean "including, but not limited to."

DETAILED DESCRIPTION OF EMBODIMENTS

Various embodiments are presented of a map tool displaying a three-dimensional view of a map, where the map tool determines whether or not partially occluded labels in the map view are to be drawn or not to be drawn in the map view.

Mobile devices and desktop computers may provide a user with map views or map navigation that includes a three-dimensional, isometric view corresponding to a current position. In some cases, the three-dimensional view may be constructed based on Global Positioning System (GPS) data, map information from other sources, or based on GPS data combined with map information from other sources. In some cases, a map view may be constructed from map information corresponding to a given address or to some other piece of information from which location information may be derived. For example, from any point on earth, a user may give a voice command to the map tool, such as "show me the front of the Metropolitan Museum of Art in New York City." In response to the voice command, the map tool, may access map information for the location of the Metropolitan Museum of Art in New York City and generate a map view to the user.

In some embodiments, the map view provides a user with an isometric view of landmarks and structures, where the landmarks and structures are represented as geometric figures such as boxes, spheres, or triangles. In some cases, the map view composed of geometric figures may be considered a low resolution proxy of the actual, or high resolution version of the surrounding environment. In either the bird's eye view,

ground level view, or the isometric view of geometric figures, the map tool may use one or more sources of map information to construct the map view. In some cases, a data source containing two-dimensional information may be combined with another data source containing three-dimensional information in order to generate a three-dimensional map view.

In an embodiment, a representation of a map area may be created, such as a grid, where the grid may be subdivided into multiple segments such that each segment is initially defined in two-dimensional (2D) space. For example, if a map area to be displayed is a 10,000 square meter area, a representation of the map area may be a grid subdivided into segments each representative of a square meter. In other examples, the segments may be defined in terms of other geometric figures or defined in terms of varying dimensions. Given the grid representation of the map area, the map tool may access three-dimensional (3D) data to extract height values for various objects within the 3D space. Upon correlating the grid area with the corresponding area in the 3D data, a height value may be determined for each grid segment. In this way, each segment may be modified to include a height value and the result is each respective grid segment of the grid area may now be defined in terms of three dimensions, the original 2D parameters of the respective grid segment and the height value from the 3D data. Given the 3D model created from extruding grid segments in the grid according to corresponding height values in the 3D data, the map tool may use the 3D model to determine whether or not a given label in a map view of the map is occluded or not based on how visible a label in the map view would be given the location and dimensions of each of the segments. In some cases, the map view is based on the 3D data or other data and not based on the generated representation of the map.

Further in this embodiment, the 3D data may be 3D mesh data for the surrounding environment. The 3D mesh data may be used to construct a model of the surrounding environment of varying degrees of granularity. While higher granularity models may result in greater computational complexity, using the higher granularity of definition for objects in the surrounding environment may also provide much greater accuracy in determining whether a given label is occluded, and how much the given label is occluded.

While in the different embodiments, different tradeoffs may be made as a result of balancing computational complexity with accuracy, the overall process of the determining label occlusion may be the similar. However, the map tool may use the granularity of the 3D information as a component in determining occlusion of a label. For example, when using highly detailed, fine grained 3D data, the map tool may determine that a single measure indicating that a label is occluded may be sufficient to make a determination not to render a label. However, when using a rough, low resolution 3D model, the map tool may determine that a higher threshold for determining whether a label is occluded must be met before making a determination not to render a label.

In an embodiment, 3D data may be 3D mesh data, which may contain data defining the location and orientation of thousands of triangles for a given map view, and where the 3D mesh data is vector data not raster data. Further in this embodiment, 2D data may be obtained from maps for a given city or county which define the locations and the dimensions of footprints for structures, roads, sidewalks, plazas, or other objects. In this embodiment, in the interest of speed and computational complexity, a 3D proxy may be constructed through the transformation of the 2D model into a 3D model using selected pieces of information from the 3D model to enhance the 2D model. For example, if the 2D model provides

information regarding the footprint of a given building, the map tool may then reference the 3D model to identify the corresponding location of the footprint of the building. Once the location of the footprint of the building is determined in the 3D model, one or more height values may be extracted from the 3D model for the building. Now, given the footprint of the building and the one or more height values, a rough box or polygon may be extruded to one of the height values, or to some value derived from the height values in order to generate an approximate 3D shape. This process may be repeated for each object in the 2D data, thereby creating a rough, low-resolution version of the surrounding environment. A benefit of using the low resolution proxy of the finer grained 3D mesh data is fewer calculations for making determinations of visibility of labels. A tradeoff may be the accuracy of a determination of occlusion of a given label.

In some embodiments, for a given map view, various areas of the map view may be labeled to provide the user additional information. A label may be an icon, a pin, a marker, an annotation, or any piece of information that may be drawn within the map view. For example, street names, park names, geographical features, and labels identifying structures or landmarks or works of art.

Detailed Description Considerations

In the following detailed description, numerous details are set forth to provide a thorough understanding of the claimed subject matter. However, it will be understood by those skilled in the art that the claimed subject matter may be practiced without these specific details. In other instances, methods, apparatus or systems that would be known by one of ordinary skill have not been described in detail so as not to obscure claimed subject matter.

It will also be understood that, although the terms first, second, etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first contact could be termed a second contact, and, similarly, a second contact could be termed a first contact, without departing from the scope of the present invention. The first contact and the second contact are both contacts, but they are not the same contact.

Embodiments of electronic devices, user interfaces for such devices, and associated processes for using such devices are described. In some embodiments, the device is a portable communications device, such as a mobile telephone, that also contains other functions, such as PDA and/or music player functions. Exemplary embodiments of portable multifunction devices include, without limitation, the iPhone®, iPod Touch®, and iPad® devices from Apple Inc. of Cupertino, Calif. Other portable electronic devices, such as laptops or tablet computers with touch-sensitive surfaces (e.g., touch screen displays and/or touch pads), may also be used. It should also be understood that, in some embodiments, the device is not a portable communications device, but is a desktop computer with a touch-sensitive surface (e.g., a touch screen display and/or a touch pad). In some embodiments, the device is a gaming computer with orientation sensors (e.g., orientation sensors in a gaming controller).

In the discussion that follows, an electronic device that includes a display and a touch-sensitive surface is described. It should be understood, however, that the electronic device may include one or more other physical user-interface devices, such as a physical keyboard, a mouse and/or a joystick.

The device typically supports a variety of applications, such as one or more of the following: a drawing application, a presentation application, a word processing application, a

website creation application, a disk authoring application, a spreadsheet application, a gaming application, a telephone application, a video conferencing application, an e-mail application, an instant messaging application, a workout support application, a photo management application, a digital camera application, a digital video camera application, a web browsing application, a digital music player application, and/or a digital video player application.

The various applications that may be executed on the device may use at least one common physical user-interface device, such as the touch-sensitive surface. One or more functions of the touch-sensitive surface as well as corresponding information displayed on the device may be adjusted and/or varied from one application to the next and/or within a respective application. In this way, a common physical architecture (such as the touch-sensitive surface) of the device may support the variety of applications with user interfaces that are intuitive and transparent to the user.

Some portions of the detailed description which follow are presented in terms of algorithms or symbolic representations of operations on binary digital signals stored within a memory of a specific apparatus or special purpose computing device or platform. In the context of this particular specification, the term specific apparatus or the like includes a general purpose computer once it is programmed to perform particular functions pursuant to instructions from program software. Algorithmic descriptions or symbolic representations are examples of techniques used by those of ordinary skill in the signal processing or related arts to convey the substance of their work to others skilled in the art. An algorithm is here, and is generally, considered to be a self-consistent sequence of operations or similar signal processing leading to a desired result. In this context, operations or processing involve physical manipulation of physical quantities. Typically, although not necessarily, such quantities may take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared or otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to such signals as bits, data, values, elements, symbols, characters, terms, numbers, numerals or the like. It should be understood, however, that all of these or similar terms are to be associated with appropriate physical quantities and are merely convenient labels. Unless specifically stated otherwise, as apparent from the following discussion, it is appreciated that throughout this specification discussions utilizing terms such as “processing”, “computing”, “calculating”, “determining”, or the like refer to actions or processes of a specific apparatus, such as a special purpose computer or a similar special purpose electronic computing device. In the context of this specification, therefore, a special purpose computer or a similar special purpose electronic computing device is capable of manipulating or transforming signals, typically represented as physical electronic or magnetic quantities within memories, registers, or other information storage devices, transmission devices, or display devices of the special purpose computer or similar special purpose electronic computing device.

Example Multifunction Device

Attention is now directed toward embodiments of portable devices with touch-sensitive displays. FIG. 1A is a block diagram illustrating portable multifunction device 100 with touch-sensitive displays 112 in accordance with some embodiments. Touch-sensitive display 112 is sometimes called a “touch screen” for convenience, and may also be known as or called a touch-sensitive display system. Device 100 may include memory 102 (which may include one or more computer-readable storage mediums, including non-

transitory computer-readable storage mediums), memory controller 122, one or more processing units (CPU’s) 120, peripherals interface 118, RF circuitry 108, audio circuitry 110, speaker 111, microphone 113, input/output (I/O) subsystem 106, other input or control devices 116, and external port 124. Device 100 may include one or more optical sensors 164. These components may communicate over one or more communication buses or signal lines 103.

It should be appreciated that device 100 is only one example of a portable multifunction device, and that device 100 may have more or fewer components than shown, may combine two or more components, or may have a different configuration or arrangement of the components. The various components shown in FIG. 1A may be implemented in hardware, software, or a combination of both hardware and software, including one or more signal processing and/or application specific integrated circuits.

Memory 102 may include high-speed random access memory and may also include non-volatile memory, such as one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices. Access to memory 102 by other components of device 100, such as CPU 120 and the peripherals interface 118, may be controlled by memory controller 122.

Peripherals interface 118 can be used to couple input and output peripherals of the device to CPU 120 and memory 102. The one or more processors 120 run or execute various software programs and/or sets of instructions stored in memory 102 to perform various functions for device 100 and to process data.

In some embodiments, peripherals interface 118, CPU 120, and memory controller 122 may be implemented on a single chip, such as chip 104. In some other embodiments, they may be implemented on separate chips.

RF (radio frequency) circuitry 108 receives and sends RF signals, also called electromagnetic signals. RF circuitry 108 converts electrical signals to/from electromagnetic signals and communicates with communications networks and other communications devices via the electromagnetic signals. RF circuitry 108 may include well-known circuitry for performing these functions, including but not limited to an antenna system, an RF transceiver, one or more amplifiers, a tuner, one or more oscillators, a digital signal processor, a CODEC chipset, a subscriber identity module (SIM) card, memory, and so forth. RF circuitry 108 may communicate with networks, such as the Internet, also referred to as the World Wide Web (WWW), an intranet and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN) and/or a metropolitan area network (MAN), and other devices by wireless communication. The wireless communication may use any of multiple communications standards, protocols and technologies, including but not limited to Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), high-speed downlink packet access (HSDPA), high-speed uplink packet access (HSUPA), wideband code division multiple access (W-CDMA), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Wireless Fidelity (Wi-Fi) (e.g., IEEE 802.11a, IEEE 802.11b, IEEE 802.11g and/or IEEE 802.11n), voice over Internet Protocol (VoIP), Wi-MAX, a protocol for e-mail (e.g., Internet message access protocol (IMAP) and/or post office protocol (POP)), instant messaging (e.g., extensible messaging and presence protocol (XMPP), Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE), Instant Messaging and Presence Service (IMPS)), and/or Short Message Service (SMS), or any other suitable communication

protocol, including communication protocols not yet developed as of the filing date of this document.

Audio circuitry **110**, speaker **111**, and microphone **113** provide an audio interface between a user and device **100**. Audio circuitry **110** receives audio data from peripherals interface **118**, converts the audio data to an electrical signal, and transmits the electrical signal to speaker **111**. Speaker **111** converts the electrical signal to human-audible sound waves. Audio circuitry **110** also receives electrical signals converted by microphone **113** from sound waves. Audio circuitry **110** converts the electrical signal to audio data and transmits the audio data to peripherals interface **118** for processing. Audio data may be retrieved from and/or transmitted to memory **102** and/or RF circuitry **108** by peripherals interface **118**. In some embodiments, audio circuitry **110** also includes a headset jack (e.g., **212**, FIG. 2). The headset jack provides an interface between audio circuitry **110** and removable audio input/output peripherals, such as output-only headphones or a headset with both output (e.g., a headphone for one or both ears) and input (e.g., a microphone).

I/O subsystem **106** couples input/output peripherals on device **100**, such as touch screen **112** and other input control devices **116**, to peripherals interface **118**. I/O subsystem **106** may include display controller **156** and one or more input controllers **160** for other input or control devices. The one or more input controllers **160** receive/send electrical signals from/to other input or control devices **116**. The other input control devices **116** may include physical buttons (e.g., push buttons, rocker buttons, etc.), dials, slider switches, joysticks, click wheels, and so forth. In some alternate embodiments, input controller(s) **160** may be coupled to any (or none) of the following: a keyboard, infrared port, USB port, and a pointer device such as a mouse. The one or more buttons (e.g., **208**, FIG. 2) may include an up/down button for volume control of speaker **111** and/or microphone **113**. The one or more buttons may include a push button (e.g., **206**, FIG. 2).

Touch-sensitive display **112** provides an input interface and an output interface between the device and a user. Display controller **156** receives and/or sends electrical signals from/to touch screen **112**. Touch screen **112** displays visual output to the user. The visual output may include graphics, text, icons, video, and any combination thereof (collectively termed “graphics”). In some embodiments, some or all of the visual output may correspond to user-interface objects.

Touch screen **112** has a touch-sensitive surface, sensor or set of sensors that accepts input from the user based on haptic and/or tactile contact. Touch screen **112** and display controller **156** (along with any associated modules and/or sets of instructions in memory **102**) detect contact (and any movement or breaking of the contact) on touch screen **112** and converts the detected contact into interaction with user-interface objects (e.g., one or more soft keys, icons, web pages or images) that are displayed on touch screen **112**. In an exemplary embodiment, a point of contact between touch screen **112** and the user corresponds to a finger of the user.

Touch screen **112** may use LCD (liquid crystal display) technology, LPD (light emitting polymer display) technology, or LED (light emitting diode) technology, although other display technologies may be used in other embodiments. Touch screen **112** and display controller **156** may detect contact and any movement or breaking thereof using any of multiple touch sensing technologies now known or later developed, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity sensor arrays or other elements for determining one or more points of contact with touch screen **112**. In an exemplary embodiment, projected mutual capacitance sens-

ing technology is used, such as that found in the iPhone®, iPod Touch®, and iPad® from Apple Inc. of Cupertino, Calif.

Touch screen **112** may have a video resolution in excess of 100 dpi. In some embodiments, the touch screen has a video resolution of approximately 160 dpi. The user may make contact with touch screen **112** using any suitable object or appendage, such as a stylus, a finger, and so forth. In some embodiments, the user interface is designed to work primarily with finger-based contacts and gestures, which can be less precise than stylus-based input due to the larger area of contact of a finger on the touch screen. In some embodiments, the device translates the rough finger-based input into a precise pointer/cursor position or command for performing the actions desired by the user.

In some embodiments, in addition to the touch screen, device **100** may include a touchpad (not shown) for activating or deactivating particular functions. In some embodiments, the touchpad is a touch-sensitive area of the device that, unlike the touch screen, does not display visual output. The touchpad may be a touch-sensitive surface that is separate from touch screen **112** or an extension of the touch-sensitive surface formed by the touch screen.

Device **100** also includes power system **162** for powering the various components. Power system **162** may include a power management system, one or more power sources (e.g., battery, alternating current (AC)), a recharging system, a power failure detection circuit, a power converter or inverter, a power status indicator (e.g., a light-emitting diode (LED)) and any other components associated with the generation, management and distribution of power in portable devices.

Device **100** may also include one or more optical sensors **164**. FIG. 1A shows an optical sensor coupled to optical sensor controller **158** in I/O subsystem **106**. Optical sensor **164** may include charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) phototransistors. Optical sensor **164** receives light from the environment, projected through one or more lens, and converts the light to data representing an image. In conjunction with imaging module **143** (also called a camera module), optical sensor **164** may capture still images or video. In some embodiments, an optical sensor is located on the back of device **100**, opposite touch screen display **112** on the front of the device, so that the touch screen display may be used as a viewfinder for still and/or video image acquisition. In some embodiments, another optical sensor is located on the front of the device so that the user’s image may be obtained for videoconferencing while the user views the other video conference participants on the touch screen display.

Device **100** may also include one or more proximity sensors **166**. FIG. 1A shows proximity sensor **166** coupled to peripherals interface **118**. Alternately, proximity sensor **166** may be coupled to input controller **160** in I/O subsystem **106**. In some embodiments, the proximity sensor turns off and disables touch screen **112** when the multifunction device is placed near the user’s ear (e.g., when the user is making a phone call).

Device **100** includes one or more orientation sensors **168**. In some embodiments, the one or more orientation sensors include one or more accelerometers (e.g., one or more linear accelerometers and/or one or more rotational accelerometers). In some embodiments, the one or more orientation sensors include one or more gyroscopes. In some embodiments, the one or more orientation sensors include one or more magnetometers. In some embodiments, the one or more orientation sensors include one or more of global positioning system (GPS), Global Navigation Satellite System (GLO-NASS), and/or other global navigation system receivers. The

GPS, GLONASS, and/or other global navigation system receivers may be used for obtaining information concerning the location and orientation (e.g., portrait or landscape) of device **100**. In some embodiments, the one or more orientation sensors include any combination of orientation/rotation sensors. FIG. **1A** shows the one or more orientation sensors **168** coupled to peripherals interface **118**. Alternately, the one or more orientation sensors **168** may be coupled to an input controller **160** in I/O subsystem **106**. In some embodiments, information is displayed on the touch screen display in a portrait view or a landscape view based on an analysis of data received from the one or more orientation sensors.

In some embodiments, the software components stored in memory **102** include operating system **126**, communication module (or set of instructions) **128**, contact/motion module (or set of instructions) **130**, graphics module (or set of instructions) **132**, text input module (or set of instructions) **134**, Global Positioning System (GPS) module (or set of instructions) **135**, and applications (or sets of instructions) **136**. Furthermore, in some embodiments memory **102** stores device/global internal state **157**, as shown in FIGS. **1A** and **3**. Device/global internal state **157** includes one or more of: active application state, indicating which applications, if any, are currently active; display state, indicating what applications, views or other information occupy various regions of touch screen display **112**; sensor state, including information obtained from the device's various sensors and input control devices **116**; and location information concerning the device's location and/or attitude.

Operating system **126** (e.g., Darwin, RTXC, LINUX, UNIX, OS X, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communication between various hardware and software components.

Communication module **128** facilitates communication with other devices over one or more external ports **124** and also includes various software components for handling data received by RF circuitry **108** and/or external port **124**. External port **124** (e.g., Universal Serial Bus (USB), FIREWIRE, etc.) is adapted for coupling directly to other devices or indirectly over a network (e.g., the Internet, wireless LAN, etc.). In some embodiments, the external port is a multi-pin (e.g., 30-pin) connector that is the same as, or similar to and/or compatible with the 30-pin connector used on iPod (trademark of Apple Inc.) devices.

Contact/motion module **130** may detect contact with touch screen **112** (in conjunction with display controller **156**) and other touch sensitive devices (e.g., a touchpad or physical click wheel). Contact/motion module **130** includes various software components for performing various operations related to detection of contact, such as determining if contact has occurred (e.g., detecting a finger-down event), determining if there is movement of the contact and tracking the movement across the touch-sensitive surface (e.g., detecting one or more finger-dragging events), and determining if the contact has ceased (e.g., detecting a finger-up event or a break in contact). Contact/motion module **130** receives contact data from the touch-sensitive surface. Determining movement of the point of contact, which is represented by a series of contact data, may include determining speed (magnitude), velocity (magnitude and direction), and/or an acceleration (a change in magnitude and/or direction) of the point of contact. These operations may be applied to single contacts (e.g., one finger contacts) or to multiple simultaneous contacts (e.g.,

“multitouch”/multiple finger contacts). In some embodiments, contact/motion module **130** and display controller **156** detect contact on a touchpad.

Contact/motion module **130** may detect a gesture input by a user. Different gestures on the touch-sensitive surface have different contact patterns. Thus, a gesture may be detected by detecting a particular contact pattern. For example, detecting a finger tap gesture includes detecting a finger-down event followed by detecting a finger-up (lift off) event at the same position (or substantially the same position) as the finger-down event (e.g., at the position of an icon). As another example, detecting a finger swipe gesture on the touch-sensitive surface includes detecting a finger-down event followed by detecting one or more finger-dragging events, and subsequently followed by detecting a finger-up (lift off) event.

Graphics module **132** includes various known software components for rendering and displaying graphics on touch screen **112** or other display, including components for changing the intensity of graphics that are displayed. As used herein, the term “graphics” includes any object that can be displayed to a user, including without limitation text, web pages, icons (such as user-interface objects including soft keys), digital images, videos, animations and the like.

In some embodiments, graphics module **132** stores data representing graphics to be used. Each graphic may be assigned a corresponding code. Graphics module **132** receives, from applications etc., one or more codes specifying graphics to be displayed along with, if necessary, coordinate data and other graphic property data, and then generates screen image data to output to display controller **156**.

Text input module **134**, which may be a component of graphics module **132**, provides soft keyboards for entering text in various applications (e.g., contacts **137**, e-mail **140**, IM **141**, browser **147**, and any other application that needs text input).

GPS module **135** determines the location of the device and provides this information for use in various applications (e.g., to telephone **138** for use in location-based dialing, to camera **143** as picture/video metadata, and to applications that provide location-based services such as weather widgets, local yellow page widgets, and map/navigation widgets).

Applications **136** may include the following modules (or sets of instructions), or a subset or superset thereof:

- contacts module **137** (sometimes called an address book or contact list);
- telephone module **138**;
- video conferencing module **139**;
- e-mail client module **140**;
- instant messaging (IM) module **141**;
- workout support module **142**;
- camera module **143** for still and/or video images;
- image management module **144**;
- browser module **147**;
- calendar module **148**;
- widget modules **149**, which may include one or more of: weather widget **149-1**, stocks widget **149-2**, calculator widget **149-3**, alarm clock widget **149-4**, dictionary widget **149-5**, and other widgets obtained by the user, as well as user-created widgets **149-6**;
- widget creator module **150** for making user-created widgets **149-6**;
- search module **151**;
- video and music player module **152**, which may be made up of a video player module and a music player module;
- notes module **153**;
- map module **154**; and/or
- online video module **155**.

11

Examples of other applications **136** that may be stored in memory **102** include other word processing applications, other image editing applications, drawing applications, presentation applications, JAVA-enabled applications, encryption, digital rights management, voice recognition, and voice replication.

In conjunction with touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, and text input module **134**, contacts module **137** may be used to manage an address book or contact list (e.g., stored in application internal state **192** of contacts module **137** in memory **102** or memory **370**), including: adding name(s) to the address book; deleting name(s) from the address book; associating telephone number(s), e-mail address(es), physical address(es) or other information with a name; associating an image with a name; categorizing and sorting names; providing telephone numbers or e-mail addresses to initiate and/or facilitate communications by telephone **138**, video conference **139**, e-mail **140**, or IM **141**; and so forth.

In conjunction with RF circuitry **108**, audio circuitry **110**, speaker **111**, microphone **113**, touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, and text input module **134**, telephone module **138** may be used to enter a sequence of characters corresponding to a telephone number, access one or more telephone numbers in address book **137**, modify a telephone number that has been entered, dial a respective telephone number, conduct a conversation and disconnect or hang up when the conversation is completed. As noted above, the wireless communication may use any of multiple communications standards, protocols and technologies.

In conjunction with RF circuitry **108**, audio circuitry **110**, speaker **111**, microphone **113**, touch screen **112**, display controller **156**, optical sensor **164**, optical sensor controller **158**, contact module **130**, graphics module **132**, text input module **134**, contact list **137**, and telephone module **138**, videoconferencing module **139** includes executable instructions to initiate, conduct, and terminate a video conference between a user and one or more other participants in accordance with user instructions.

In conjunction with RF circuitry **108**, touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, and text input module **134**, e-mail client module **140** includes executable instructions to create, send, receive, and manage e-mail in response to user instructions. In conjunction with image management module **144**, e-mail client module **140** makes it very easy to create and send e-mails with still or video images taken with camera module **143**.

In conjunction with RF circuitry **108**, touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, and text input module **134**, the instant messaging module **141** includes executable instructions to enter a sequence of characters corresponding to an instant message, to modify previously entered characters, to transmit a respective instant message (for example, using a Short Message Service (SMS) or Multimedia Message Service (MMS) protocol for telephony-based instant messages or using XMPP, SIMPLE, or IMPS for Internet-based instant messages), to receive instant messages and to view received instant messages. In some embodiments, transmitted and/or received instant messages may include graphics, photos, audio files, video files and/or other attachments as are supported in a MMS and/or an Enhanced Messaging Service (EMS). As used herein, "instant messaging" refers to both telephony-based messages (e.g., messages sent using SMS or MMS) and Internet-based messages (e.g., messages sent using XMPP, SIMPLE, or IMPS).

12

In conjunction with RF circuitry **108**, touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, text input module **134**, GPS module **135**, map module **154**, and music player module **146**, workout support module **142** includes executable instructions to create workouts (e.g., with time, distance, and/or calorie burning goals); communicate with workout sensors (sports devices); receive workout sensor data; calibrate sensors used to monitor a workout; select and play music for a workout; and display, store and transmit workout data.

In conjunction with touch screen **112**, display controller **156**, optical sensor(s) **164**, optical sensor controller **158**, contact module **130**, graphics module **132**, and image management module **144**, camera module **143** includes executable instructions to capture still images or video (including a video stream) and store them into memory **102**, modify characteristics of a still image or video, or delete a still image or video from memory **102**.

In conjunction with touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, text input module **134**, and camera module **143**, image management module **144** includes executable instructions to arrange, modify (e.g., edit), or otherwise manipulate, label, delete, present (e.g., in a digital slide show or album), and store still and/or video images.

In conjunction with RF circuitry **108**, touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, and text input module **134**, browser module **147** includes executable instructions to browse the Internet in accordance with user instructions, including searching, linking to, receiving, and displaying web pages or portions thereof, as well as attachments and other files linked to web pages.

In conjunction with RF circuitry **108**, touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, text input module **134**, e-mail client module **140**, and browser module **147**, calendar module **148** includes executable instructions to create, display, modify, and store calendars and data associated with calendars (e.g., calendar entries, to do lists, etc.) in accordance with user instructions.

In conjunction with RF circuitry **108**, touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, text input module **134**, and browser module **147**, widget modules **149** are mini-applications that may be downloaded and used by a user (e.g., weather widget **149-1**, stocks widget **149-2**, calculator widget **149-3**, alarm clock widget **149-4**, and dictionary widget **149-5**) or created by the user (e.g., user-created widget **149-6**). In some embodiments, a widget includes an HTML (Hypertext Markup Language) file, a CSS (Cascading Style Sheets) file, and a JavaScript file. In some embodiments, a widget includes an XML (Extensible Markup Language) file and a JavaScript file (e.g., Yahoo! Widgets).

In conjunction with RF circuitry **108**, touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, text input module **134**, and browser module **147**, the widget creator module **150** may be used by a user to create widgets (e.g., turning a user-specified portion of a web page into a widget).

In conjunction with touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, and text input module **134**, search module **151** includes executable instructions to search for text, music, sound, image, video, and/or other files in memory **102** that match one or more search criteria (e.g., one or more user-specified search terms) in accordance with user instructions.

In conjunction with touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, audio circuitry **110**, speaker **111**, RF circuitry **108**, and browser module **147**, video and music player module **152** includes executable instructions that allow the user to download and play back recorded music and other sound files stored in one or more file formats, such as MP3 or AAC files, and executable instructions to display, present or otherwise play back videos (e.g., on touch screen **112** or on an external, connected display via external port **124**). In some embodiments, device **100** may include the functionality of an MP3 player, such as an iPod (trademark of Apple Inc.).

In conjunction with touch screen **112**, display controller **156**, contact module **130**, graphics module **132**, and text input module **134**, notes module **153** includes executable instructions to create and manage notes, to do lists, and the like in accordance with user instructions.

In conjunction with RF circuitry **108**, touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, text input module **134**, GPS module **135**, and browser module **147**, map module **154** may be used to receive, display, modify, and store maps and data associated with maps (e.g., driving directions; data on stores and other points of interest at or near a particular location; and other location-based data) in accordance with user instructions.

In conjunction with touch screen **112**, display system controller **156**, contact module **130**, graphics module **132**, audio circuitry **110**, speaker **111**, RF circuitry **108**, text input module **134**, e-mail client module **140**, and browser module **147**, online video module **155** includes instructions that allow the user to access, browse, receive (e.g., by streaming and/or download), play back (e.g., on the touch screen or on an external, connected display via external port **124**), send an e-mail with a link to a particular online video, and otherwise manage online videos in one or more file formats, such as H.264. In some embodiments, instant messaging module **141**, rather than e-mail client module **140**, is used to send a link to a particular online video.

Each of the above identified modules and applications correspond to a set of executable instructions for performing one or more functions described above and the methods described in this application (e.g., the computer-implemented methods and other information processing methods described herein). These modules (i.e., sets of instructions) need not be implemented as separate software programs, procedures or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various embodiments. In some embodiments, memory **102** may store a subset of the modules and data structures identified above. Furthermore, memory **102** may store additional modules and data structures not described above.

In some embodiments, device **100** is a device where operation of a predefined set of functions on the device is performed exclusively through a touch screen and/or a touchpad. By using a touch screen and/or a touchpad as the primary input control device for operation of device **100**, the number of physical input control devices (such as push buttons, dials, and the like) on device **100** may be reduced.

The predefined set of functions that may be performed exclusively through a touch screen and/or a touchpad include navigation between user interfaces. In some embodiments, the touchpad, when touched by the user, navigates device **100** to a main, home, or root menu from any user interface that may be displayed on device **100**. In such embodiments, the touchpad may be referred to as a "menu button." In some

other embodiments, the menu button may be a physical push button or other physical input control device instead of a touchpad.

FIG. 1B is a block diagram illustrating exemplary components for event handling in accordance with some embodiments. In some embodiments, memory **102** (in FIG. 1A) or **370** (FIG. 3) includes event sorter **170** (e.g., in operating system **126**) and a respective application **136-1** (e.g., any of the aforementioned applications **137-151**, **155**, **380-390**).

Event sorter **170** receives event information and determines the application **136-1** and application view **191** of application **136-1** to which to deliver the event information. Event sorter **170** includes event monitor **171** and event dispatcher module **174**. In some embodiments, application **136-1** includes application internal state **192**, which indicates the current application view(s) displayed on touch sensitive display **112** when the application is active or executing. In some embodiments, device/global internal state **157** is used by event sorter **170** to determine which application(s) is (are) currently active, and application internal state **192** is used by event sorter **170** to determine application views **191** to which to deliver event information.

In some embodiments, application internal state **192** includes additional information, such as one or more of: resume information to be used when application **136-1** resumes execution, user interface state information that indicates information being displayed or that is ready for display by application **136-1**, a state queue for enabling the user to go back to a prior state or view of application **136-1**, and a redo/undo queue of previous actions taken by the user.

Event monitor **171** receives event information from peripherals interface **118**. Event information includes information about a sub-event (e.g., a user touch on touch sensitive display **112**, as part of a multi-touch gesture). Peripherals interface **118** transmits information it receives from I/O subsystem **106** or a sensor, such as proximity sensor **166**, orientation sensor(s) **168**, and/or microphone **113** (through audio circuitry **110**). Information that peripherals interface **118** receives from I/O subsystem **106** includes information from touch-sensitive display **112** or a touch-sensitive surface.

In some embodiments, event monitor **171** sends requests to the peripherals interface **118** at predetermined intervals. In response, peripherals interface **118** transmits event information. In other embodiments, peripheral interface **118** transmits event information only when there is a significant event (e.g., receiving an input above a predetermined noise threshold and/or for more than a predetermined duration).

In some embodiments, event sorter **170** also includes a hit view determination module **172** and/or an active event recognizer determination module **173**.

Hit view determination module **172** provides software procedures for determining where a sub-event has taken place within one or more views, when touch sensitive display **112** displays more than one view. Views are made up of controls and other elements that a user can see on the display.

Another aspect of the user interface associated with an application is a set of views, sometimes herein called application views or user interface windows, in which information is displayed and touch-based gestures occur. The application views (of a respective application) in which a touch is detected may correspond to programmatic levels within a programmatic or view hierarchy of the application. For example, the lowest level view in which a touch is detected may be called the hit view, and the set of events that are recognized as proper inputs may be determined based, at least in part, on the hit view of the initial touch that begins a touch-based gesture.

15

Hit view determination module **172** receives information related to sub-events of a touch-based gesture. When an application has multiple views organized in a hierarchy, hit view determination module **172** identifies a hit view as the lowest view in the hierarchy which should handle the sub-event. In most circumstances, the hit view is the lowest level view in which an initiating sub-event occurs (i.e., the first sub-event in the sequence of sub-events that form an event or potential event). Once the hit view is identified by the hit view determination module, the hit view typically receives all sub-events related to the same touch or input source for which it was identified as the hit view.

Active event recognizer determination module **173** determines which view or views within a view hierarchy should receive a particular sequence of sub-events. In some embodiments, active event recognizer determination module **173** determines that only the hit view should receive a particular sequence of sub-events. In other embodiments, active event recognizer determination module **173** determines that all views that include the physical location of a sub-event are actively involved views, and therefore determines that all actively involved views should receive a particular sequence of sub-events. In other embodiments, even if touch sub-events were entirely confined to the area associated with one particular view, views higher in the hierarchy would still remain as actively involved views.

Event dispatcher module **174** dispatches the event information to an event recognizer (e.g., event recognizer **180**). In embodiments including active event recognizer determination module **173**, event dispatcher module **174** delivers the event information to an event recognizer determined by active event recognizer determination module **173**. In some embodiments, event dispatcher module **174** stores in an event queue the event information, which is retrieved by a respective event receiver module **182**.

In some embodiments, operating system **126** includes event sorter **170**. Alternatively, application **136-1** includes event sorter **170**. In yet other embodiments, event sorter **170** is a stand-alone module, or a part of another module stored in memory **102**, such as contact/motion module **130**.

In some embodiments, application **136-1** includes multiple event handlers **190** and one or more application views **191**, each of which includes instructions for handling touch events that occur within a respective view of the application's user interface. Each application view **191** of the application **136-1** includes one or more event recognizers **180**. Typically, a respective application view **191** includes multiple event recognizers **180**. In other embodiments, one or more of event recognizers **180** are part of a separate module, such as a user interface kit (not shown) or a higher level object from which application **136-1** inherits methods and other properties. In some embodiments, a respective event handler **190** includes one or more of: data updater **176**, object updater **177**, GUI updater **178**, and/or event data **179** received from event sorter **170**. Event handler **190** may utilize or call data updater **176**, object updater **177** or GUI updater **178** to update the application internal state **192**. Alternatively, one or more of the application views **191** includes one or more respective event handlers **190**. Also, in some embodiments, one or more of data updater **176**, object updater **177**, and GUI updater **178** are included in a respective application view **191**.

A respective event recognizer **180** receives event information (e.g., event data **179**) from event sorter **170**, and identifies an event from the event information. Event recognizer **180** includes event receiver **182** and event comparator **184**. In some embodiments, event recognizer **180** also includes at

16

least a subset of: metadata **183**, and event delivery instructions **188** (which may include sub-event delivery instructions).

Event receiver **182** receives event information from event sorter **170**. The event information includes information about a sub-event, for example, a touch or a touch movement. Depending on the sub-event, the event information also includes additional information, such as location of the sub-event. When the sub-event concerns motion of a touch the event information may also include speed and direction of the sub-event. In some embodiments, events include rotation of the device from one orientation to another (e.g., from a portrait orientation to a landscape orientation, or vice versa), and the event information includes corresponding information about the current orientation (also called device attitude) of the device.

Event comparator **184** compares the event information to predefined event or sub-event definitions and, based on the comparison, determines an event or sub-event, or determines or updates the state of an event or sub-event. In some embodiments, event comparator **184** includes event definitions **186**. Event definitions **186** contain definitions of events (e.g., predefined sequences of sub-events), for example, event **1** (**187-1**), event **2** (**187-2**), and others. In some embodiments, sub-events in an event **187** include, for example, touch begin, touch end, touch movement, touch cancellation, and multiple touching. In one example, the definition for event **1** (**187-1**) is a double tap on a displayed object. The double tap, for example, includes a first touch (touch begin) on the displayed object for a predetermined phase, a first lift-off (touch end) for a predetermined phase, a second touch (touch begin) on the displayed object for a predetermined phase, and a second lift-off (touch end) for a predetermined phase. In another example, the definition for event **2** (**187-2**) is a dragging on a displayed object. The dragging, for example, includes a touch (or contact) on the displayed object for a predetermined phase, a movement of the touch across touch-sensitive display **112**, and lift-off of the touch (touch end). In some embodiments, the event also includes information for one or more associated event handlers **190**.

In some embodiments, event definition **187** includes a definition of an event for a respective user-interface object. In some embodiments, event comparator **184** performs a hit test to determine which user-interface object is associated with a sub-event. For example, in an application view in which three user-interface objects are displayed on touch-sensitive display **112**, when a touch is detected on touch-sensitive display **112**, event comparator **184** performs a hit test to determine which of the three user-interface objects is associated with the touch (sub-event). If each displayed object is associated with a respective event handler **190**, the event comparator uses the result of the hit test to determine which event handler **190** should be activated. For example, event comparator **184** selects an event handler associated with the sub-event and the object triggering the hit test.

In some embodiments, the definition for a respective event **187** also includes delayed actions that delay delivery of the event information until after it has been determined whether the sequence of sub-events does or does not correspond to the event recognizer's event type.

When a respective event recognizer **180** determines that the series of sub-events do not match any of the events in event definitions **186**, the respective event recognizer **180** enters an event impossible, event failed, or event ended state, after which it disregards subsequent sub-events of the touch-based gesture. In this situation, other event recognizers, if any, that

remain active for the hit view continue to track and process sub-events of an ongoing touch-based gesture.

In some embodiments, a respective event recognizer **180** includes metadata **183** with configurable properties, flags, and/or lists that indicate how the event delivery system should perform sub-event delivery to actively involved event recognizers. In some embodiments, metadata **183** includes configurable properties, flags, and/or lists that indicate how event recognizers may interact with one another. In some embodiments, metadata **183** includes configurable properties, flags, and/or lists that indicate whether sub-events are delivered to varying levels in the view or programmatic hierarchy.

In some embodiments, a respective event recognizer **180** activates event handler **190** associated with an event when one or more particular sub-events of an event are recognized. In some embodiments, a respective event recognizer **180** delivers event information associated with the event to event handler **190**. Activating an event handler **190** is distinct from sending (and deferred sending) sub-events to a respective hit view. In some embodiments, event recognizer **180** throws a flag associated with the recognized event, and event handler **190** associated with the flag catches the flag and performs a predefined process.

In some embodiments, event delivery instructions **188** include sub-event delivery instructions that deliver event information about a sub-event without activating an event handler. Instead, the sub-event delivery instructions deliver event information to event handlers associated with the series of sub-events or to actively involved views. Event handlers associated with the series of sub-events or with actively involved views receive the event information and perform a predetermined process.

In some embodiments, data updater **176** creates and updates data used in application **136-1**. For example, data updater **176** updates the telephone number used in contacts module **137**, or stores a video file used in video player module **145**. In some embodiments, object updater **177** creates and updates objects used in application **136-1**. For example, object updater **176** creates a new user-interface object or updates the position of a user-interface object. GUI updater **178** updates the GUI. For example, GUI updater **178** prepares display information and sends it to graphics module **132** for display on a touch-sensitive display.

In some embodiments, event handler(s) **190** includes or has access to data updater **176**, object updater **177**, and GUI updater **178**. In some embodiments, data updater **176**, object updater **177**, and GUI updater **178** are included in a single module of a respective application **136-1** or application view **191**. In other embodiments, they are included in two or more software modules.

It shall be understood that the foregoing discussion regarding event handling of user touches on touch-sensitive displays also applies to other forms of user inputs to operate multifunction devices **100** with input-devices, not all of which are initiated on touch screens, e.g., coordinating mouse movement and mouse button presses with or without single or multiple keyboard presses or holds, user movements taps, drags, scrolls, etc., on touch-pads, pen stylus inputs, movement of the device, oral instructions, detected eye movements, biometric inputs, and/or any combination thereof, which may be utilized as inputs corresponding to sub-events which define an event to be recognized.

FIG. **2** illustrates a portable multifunction device **100** having a touch screen **112** in accordance with some embodiments. The touch screen may display one or more graphics within user interface (UI) **200**. In this embodiment, as well as others described below, a user may select one or more of the

graphics by making a gesture on the graphics, for example, with one or more fingers **202** (not drawn to scale in the figure) or one or more styluses **203** (not drawn to scale in the figure). In some embodiments, selection of one or more graphics occurs when the user breaks contact with the one or more graphics. In some embodiments, the gesture may include one or more taps, one or more swipes (from left to right, right to left, upward and/or downward) and/or a rolling of a finger (from right to left, left to right, upward and/or downward) that has made contact with device **100**. In some embodiments, inadvertent contact with a graphic may not select the graphic. For example, a swipe gesture that sweeps over an application icon may not select the corresponding application when the gesture corresponding to selection is a tap.

Device **100** may also include one or more physical buttons, such as “home” or menu button **204**. As described previously, menu button **204** may be used to navigate to any application **136** in a set of applications that may be executed on device **100**. Alternatively, in some embodiments, the menu button is implemented as a soft key in a GUI displayed on touch screen **112**.

In one embodiment, device **100** includes touch screen **112**, menu button **204**, push button **206** for powering the device on/off and locking the device, volume adjustment button(s) **208**, Subscriber Identity Module (SIM) card slot **210**, head set jack **212**, and docking/charging external port **124**. Push button **206** may be used to turn the power on/off on the device by depressing the button and holding the button in the depressed state for a predefined time interval; to lock the device by depressing the button and releasing the button before the predefined time interval has elapsed; and/or to unlock the device or initiate an unlock process. In an alternative embodiment, device **100** also may accept verbal input for activation or deactivation of some functions through microphone **113**.

It should be noted that, although many of the following examples will be given with reference to inputs on touch screen **112** (where the touch sensitive surface and the display are combined), a touch-sensitive surface that is separate from the display may be used instead of touch screen **112**.

Map Service Operating Environment

Various embodiments of a map tool may operate within a map service operating environment. FIG. **9** illustrates a map service operating environment, according to some embodiments. A map service **930** may provide map services for one or more client devices **902a-902c** in communication with the map service **930** through various communication methods and protocols. A map service **930** generally may provide map information and other map-related data, such as two-dimensional map image data (e.g., aerial view of roads utilizing satellite imagery), three-dimensional map image data (e.g., traversable map with three-dimensional features, such as buildings), route and direction calculation (e.g., ferry route calculations or directions between two points for a pedestrian), real-time navigation data (e.g., turn-by-turn visual navigation data in two or three dimensions), location data (e.g., where is the client device currently located), and other geographic data (e.g., wireless network coverage, weather, traffic information, or nearby points-of-interest). In various embodiments, the map service data may include localized labels for different countries or regions; localized labels may be utilized to present map labels (e.g., street names, city names, points of interest) in different languages on client devices. Client devices **902a-902c** may utilize these map services by obtaining map service data. Client devices **902a-902c** may implement various techniques to process map service data. Client devices **902a-902c** may then provide map services to various entities, including, but not limited to,

users, internal software or hardware modules, and/or other systems or devices external to the client devices **902a-902c**.

In some embodiments, a map service may be implemented by one or more nodes in a distributed computing system. Each node may be assigned one or more services or components of a map service. Some nodes may be assigned the same map service or component of a map service. A load balancing node may distribute access or requests to other nodes within a map service. In some embodiments a map service may be implemented as a single system, such as a single server. Different modules or hardware devices within a server may implement one or more of the various services provided by a map service.

A map service may provide map services by generating map service data in various formats. In some embodiments, one format of map service data may be map image data. Map image data may provide image data to a client device so that the client device may process the image data (e.g., rendering and/or displaying the image data as a two-dimensional or three-dimensional map). Map image data, whether in two or three dimensions, may specify one or more map tiles. A map tile may be a portion of a larger map image. Assembling together the map tiles of a map may produce the original map. Tiles may be generated from map image data, routing or navigation data, or any other map service data. In some embodiments map tiles may be raster-based map tiles, with tile sizes ranging from any size both larger and smaller than a commonly-used 256 pixel by 256 pixel tile. Raster-based map tiles may be encoded in any number of standard digital image representations including, but not limited to, Bitmap (.bmp), Graphics Interchange Format (.gif), Joint Photographic Experts Group (.jpg, .jpeg, etc.), Portable Networks Graphic (.png), or Tagged Image File Format (.tiff). In some embodiments, map tiles may be vector-based map tiles, encoded using vector graphics, including, but not limited to, Scalable Vector Graphics (.svg) or a Drawing File (.drw). Embodiments may also include tiles with a combination of vector and raster data. Metadata or other information pertaining to the map tile may also be included within or along with a map tile, providing further map service data to a client device. In various embodiments, a map tile may be encoded for transport utilizing various standards and/or protocols, some of which are described in examples below.

In various embodiments, map tiles may be constructed from image data of different resolutions depending on zoom level. For instance, for low zoom level (e.g., world or globe view), the resolution of map or image data need not be as high relative to the resolution at a high zoom level (e.g., city or street level). For example, when in a globe view, there may be no need to render street level artifacts as such objects would be so small as to be negligible in many cases.

A map service may perform various techniques to analyze a map tile before encoding the tile for transport. This analysis may optimize map service performance for both client devices and a map service. In some embodiments map tiles may be analyzed for complexity, according to vector-based graphic techniques, and constructed utilizing complex and non-complex layers. Map tiles may also be analyzed for common image data or patterns that may be rendered as image textures and constructed by relying on image masks. In some embodiments, raster-based image data in a map tile may contain certain mask values, which are associated with one or more textures. Embodiments may also analyze map tiles for specified features that may be associated with certain map styles that contain style identifiers.

Other map services may generate map service data relying upon various data formats separate from a map tile. For example, map services that provide location data may utilize

data formats conforming to location service protocols, such as, but not limited to, Radio Resource Location services Protocol (RRLP), TIA 801 for Code Division Multiple Access (CDMA), Radio Resource Control (RRC) position protocol, or LTE Positioning Protocol (LPP). Embodiments may also receive or request data from client devices identifying device capabilities or attributes (e.g., hardware specifications or operating system version) or communication capabilities (e.g., device communication bandwidth as determined by wireless signal strength or wire or wireless network type).

A map service may obtain map service data from internal or external sources. For example, satellite imagery used in map image data may be obtained from external services, or internal systems, storage devices, or nodes. Other examples may include, but are not limited to, GPS assistance servers, wireless network coverage databases, business or personal directories, weather data, government information (e.g., construction updates or road name changes), or traffic reports. Some embodiments of a map service may update map service data (e.g., wireless network coverage) for analyzing future requests from client devices.

Various embodiments of a map service may respond to client device requests for map services. These requests may be a request for a specific map or portion of a map. Embodiments may format requests for a map as requests for certain map tiles. In some embodiments, requests may also supply the map service with starting locations (or current locations) and destination locations for a route calculation. A client device may also request map service rendering information, such as map textures or stylesheets. In at least some embodiments, requests may also be one of a series of requests implementing turn-by-turn navigation. Requests for other geographic data may include, but are not limited to, current location, wireless network coverage, weather, traffic information, or nearby points-of-interest.

A map service may, in some embodiments, may analyze client device requests to optimize a device or map service operation. For example, a map service may recognize that the location of a client device is in an area of poor communications (e.g., weak wireless signal) and send more map service data to supply a client device in the event of loss in communication or send instructions to utilize different client hardware (e.g., orientation sensors) or software (e.g., utilize wireless location services or Wi-Fi positioning instead of GPS-based services). In another example, a map service may analyze a client device request for vector-based map image data and determine that raster-based map data better optimizes the map image data according to the image's complexity. Embodiments of other map services may perform similar analysis on client device requests and as such the above examples are not intended to be limiting.

Various embodiments of client devices (e.g., client devices **902a-902c**) may be implemented on different device types. Examples of a portable-multifunction device include the devices illustrated in FIGS. 1 through 3 and 9, such as multifunction device **100** and multifunction device **300**. Client devices **902a-902c** may utilize map service **930** through various communication methods and protocols described below. In some embodiments, client devices **902a-902c** may obtain map service data from map service **930**. Client devices **902a-902c** may request or receive map service data. Client devices **902a-902c** may then process map service data (e.g., render and/or display the data) and may send the data to another software or hardware module on the device or to an external device or system.

A client device may, according to some embodiments, implement techniques to render and/or display maps. These

maps may be requested or received in various formats, such as map tiles described above. A client device may render a map in two-dimensional or three-dimensional views. Some embodiments of a client device may display a rendered map and allow a user, system, or device providing input to manipulate a virtual camera in the map, changing the map display according to the virtual camera's position, orientation, and field-of-view. Various forms and input devices may be implemented to manipulate a virtual camera. In some embodiments, touch input, through certain single or combination gestures (e.g., touch-and-hold or a swipe) may manipulate the virtual camera. Other embodiments may allow manipulation of the device's physical location to manipulate a virtual camera. For example, a client device may be tilted up from its current position to manipulate the virtual camera to rotate up. In another example, a client device may be tilted forward from its current position to move the virtual camera forward. Other input devices to the client device may be implemented including, but not limited to, auditory input (e.g., spoken words), a physical keyboard, mouse, and/or a joystick.

Embodiments may provide various visual feedback to virtual camera manipulations, such as displaying an animation of possible virtual camera manipulations when transitioning from two-dimensional map views to three-dimensional map views. Embodiments may also allow input to select a map feature or object (e.g., a building) and highlight the object, producing a blur effect that maintains the virtual camera's perception of three-dimensional space.

In some embodiments, a client device may implement a navigation system (e.g., turn-by-turn navigation). A navigation system provides directions or route information, which may be displayed to a user. Embodiments of a client device may request directions or a route calculation from a map service. A client device may receive map image data and route data from a map service. In some embodiments, a client device may implement a turn-by-turn navigation system, which provides real-time route and direction information based upon location information and route information received from a map service and/or other location system, such as Global Positioning Satellite (GPS). A client device may display map image data that reflects the current location of the client device and update the map image data in real-time. A navigation system may provide auditory or visual directions to follow a certain route.

A virtual camera may be implemented to manipulate navigation map data according to some embodiments. Some embodiments of client devices may allow the device to adjust the virtual camera display orientation to bias toward the route destination. Embodiments may also allow virtual camera to navigation turns simulating the inertial motion of the virtual camera.

Client devices may implement various techniques to utilize map service data from map service. Embodiments may implement some techniques to optimize rendering of two-dimensional and three-dimensional map image data. In some embodiments, a client device may locally store rendering information. For example, a client may store a stylesheet which provides rendering directions for image data containing style identifiers. In another example, common image textures may be stored to decrease the amount of map image data transferred from a map service. Client devices may also implement various modeling techniques to render two-dimensional and three-dimensional map image data, examples of which include, but are not limited to: generating three-dimensional buildings out of two-dimensional building footprint data; modeling two-dimensional and three-dimensional map objects to determine the client device communication

environment; generating models to determine whether map labels are seen from a certain virtual camera position; and generating models to smooth transitions between map image data. Some embodiments of client devices may also order or prioritize map service data in certain techniques. For example, a client device may detect the motion or velocity of a virtual camera, which if exceeding certain threshold values, lower-detail image data will be loaded and rendered of certain areas. Other examples include: rendering vector-based curves as a series of points, preloading map image data for areas of poor communication with a map service, adapting textures based on display zoom level, or rendering map image data according to complexity.

In some embodiments, client devices may communicate utilizing various data formats separate from a map tile. For example, some client devices may implement Assisted Global Positioning Satellites (A-GPS) and communicate with location services that utilize data formats conforming to location service protocols, such as, but not limited to, Radio Resource Location services Protocol (RRLP), TIA 801 for Code Division Multiple Access (CDMA), Radio Resource Control (RRC) position protocol, or LTE Positioning Protocol (LPP). Client devices may also receive GPS signals directly. Embodiments may also send data, with or without solicitation from a map service, identifying the client device's capabilities or attributes (e.g., hardware specifications or operating system version) or communication capabilities (e.g., device communication bandwidth as determined by wireless signal strength or wire or wireless network type).

FIG. 9 illustrates one possible embodiment of an operating environment 900 for a map service 930 and client devices 902a-902c. In some embodiments, devices 902a, 902b, and 902c can communicate over one or more wire or wireless networks 910. For example, wireless network 910, such as a cellular network, can communicate with a wide area network (WAN) 920, such as the Internet, by use of gateway 914. A gateway 914 may provide a packet oriented mobile data service, such as General Packet Radio Service (GPRS), or other mobile data service allowing wireless networks to transmit data to other networks, such as wide area network 920. Likewise, access device 912 (e.g., IEEE 802.11g wireless access device) can provide communication access to WAN 920. Devices 902a and 902b can be any portable electronic or computing device capable of communicating with a map service, such as a portable multifunction device described below with respect to FIGS. 1 to 3 and 9. Device 902c can be any non-portable electronic or computing device capable of communicating with a map service, such as a system described below in FIG. 9.

In some embodiments, both voice and data communications can be established over wireless network 910 and access device 912. For example, device 902a can place and receive phone calls (e.g., using voice over Internet Protocol (VoIP) protocols), send and receive e-mail messages (e.g., using Simple Mail Transfer Protocol (SMTP) or Post Office Protocol 3 (POP3)), and retrieve electronic documents and/or streams, such as web pages, photographs, and videos, over wireless network 910, gateway 914, and WAN 920 (e.g., using Transmission Control Protocol/Internet Protocol (TCP/IP) or User Datagram Protocol (UDP)). Likewise, in some implementations, devices 902b and 902c can place and receive phone calls, send and receive e-mail messages, and retrieve electronic documents over access device 912 and WAN 920. In various embodiments, any of the illustrated client device may communicate with map service 930 and/or other service(s) 950 using a persistent connection established

in accordance with one or more security protocols, such as the Secure Sockets Layer (SSL) protocol or the Transport Layer Security (TLS) protocol.

Devices **902a** and **902b** can also establish communications by other means. For example, wireless device **902a** can communicate with other wireless devices (e.g., other devices **902a** or **902b**, cell phones) over the wireless network **910**. Likewise devices **902a** and **902b** can establish peer-to-peer communications **942** (e.g., a personal area network) by use of one or more communication subsystems, such as Bluetooth® communication from Bluetooth Special Interest Group, Inc. of Kirkland, Wash. **902c** can also establish peer to peer communications with devices **902a** or **902b**. (not pictured). Other communication protocols and topologies can also be implemented. Devices **902a** and **902b** may also receive Global Positioning Satellite (GPS) signals from GPS **940**.

Devices **902a**, **902b**, and **902c** can communicate with map service **930** over the one or more wire and/or wireless networks, **910** or **912**. For example, map service **930** can provide a map service data to rendering devices **902a**, **902b**, and **902c**. Map service **930** may also communicate with other services **950** to obtain data to implement map services. Map service **930** and other services **950** may also receive GPS signals from GPS **940**.

In various embodiments, map service **930** and/or other service(s) **950** may be configured to process search requests from any of client devices. Search requests may include but are not limited to queries for business, address, residential locations, points of interest, or some combination thereof. Map service **930** and/or other service(s) **950** may be configured to return results related to a variety of parameters including but not limited to a location entered into an address bar or other text entry field (including abbreviations and/or other shorthand notation), a current map view (e.g., user may be viewing one location on the multifunction device while residing in another location), current location of the user (e.g., in cases where the current map view did not include search results), and the current route (if any). In various embodiments, these parameters may affect the composition of the search results (and/or the ordering of the search results) based on different priority weightings. In various embodiments, the search results that are returned may be a subset of results selected based on specific criteria include but not limited to a quantity of times the search result (e.g., a particular point of interest) has been requested, a measure of quality associated with the search result (e.g., highest user or editorial review rating), and/or the volume of reviews for the search results (e.g., the number of times the search result has been review or rated).

In various embodiments, map service **930** and/or other service(s) **950** may be configured to provide auto-complete search results that may be displayed on the client device, such as within the mapping application. For instance, auto-complete search results may populate a portion of the screen as the user enters one or more search keywords on the multifunction device. In some cases, this feature may save the user time as the desired search result may be displayed before the user enters the full search query. In various embodiments, the auto complete search results may be search results found by the client on the client device (e.g., bookmarks or contacts), search results found elsewhere (e.g., from the internet) by map service **930** and/or other service(s) **950**, and/or some combination thereof. As is the case with commands, any of the search queries may be entered by the user via voice or through typing. The multifunction device may be configured to display search results graphically within any of the map display described herein. For instance, a pin or other graphi-

cal indicator may specify locations of search results as points of interest. In various embodiments, responsive to a user selection of one of these points of interest (e.g., a touch selection, such as a tap), the multifunction device may be configured to display additional information about the selected point of interest including but not limited to ratings, reviews or review snippets, hours of operation, store status (e.g., open for business, permanently closed, etc.), and/or images of a storefront for the point of interest. In various embodiments, any of this information may be displayed on a graphical information card that is displayed in response to the user's selection of the point of interest.

In various embodiments, map service **930** and/or other service(s) **950** may provide one or more feedback mechanisms to receive feedback from client devices **902a-c**. For instance, client devices may provide feedback on search results to map service **930** and/or other service(s) **950** (e.g., feedback specifying ratings, reviews, temporary or permanent business closures, errors etc.); this feedback may be used to update information about points of interest in order to provide more accurate or more up-to-date search results in the future. In some embodiments, map service **930** and/or other service(s) **950** may provide testing information to the client device (e.g., an AB test) to determine which search results are best. For instance, at random intervals, the client device may receive and present two search results to a user and allow the user to indicate the best result. The client device may report the test results to map service **930** and/or other service(s) **950** to improve future search results based on the chosen testing technique, such as an A/B test technique in which a baseline control sample is compared to a variety of single-variable test samples in order to improve results.

Example Mapping Functionality

FIG. 3 illustrates another example of a multifunction device that may implement a map tool in accord with the embodiments described, where the multifunction device may be configured in a manner similar to the multifunction device described above. In the illustrated embodiment, a multifunction device **300** includes a mapping application (e.g., map module **154** described above) that may be stored in one or more memories of multifunction device **300** and executed on one or more processors of multifunction device **300**. As is the case for the multifunction device described above, multifunction device **300** may include one or more controls **302** for operating the multifunction device. These controls may include but are not limited to power controls for turning the device on and off, volume controls for adjusting the ear piece volume or the speaker volume, menu controls for navigation functions of the device, and function controls for initiating one or more function or actions on the device. Controls **302** may include hardware controls or software controls. For instance, the bottom left corner of electronic display **312** includes a graphical representation of a control **312** that may be selected by a user, such as by way of touch in accordance with the touch screen functionality described above.

Multifunction device **300** may also include other components similar to those described above, such as a microphone **304**, an earpiece **306** (e.g., a speaker through which to convey audio representations of telephone calls), an optical sensor **308**, and/or a speaker **310**. Each of these components may be configured in a similar manner to those like-named components of FIG. 2 described above. Furthermore, electronic display **312** may be configured with touch screen capability, such as touch screen **112** described above. In various embodiments, controls (e.g., on screen control(s) **302**) may be utilized to perform any of a variety of map-related functions including but not limited to zoom in, zoom out, rotate screen,

pan screen, toggle views (e.g., two-dimensions to three dimensions and vice versa), and/or another map related activity. In various embodiments, one or more gestures may be utilized to perform any of the aforesaid map controls (with or without the use of an actual graphical on-screen control). In one non-limiting example, a one figure gesture may be utilized to adjust the pitch within a three-dimensional map view.

As noted above, multifunction device **300** includes a mapping application that may be stored in one or more memories of multifunction device **300** and executed on one or more processors of multifunction device **300**. In the illustrated embodiment, the graphical representation of the mapping application may include a map **314** of a geographic region. This map may be presented as a two-dimensional map or a three-dimensional map, the selection of which may be specified through, e.g., a user-configurable parameter of the mapping application. In some embodiments, the multifunction device may toggle between two-dimensional map or three-dimensional map views responsive to input from any input component of the multifunction device. In one non-limiting example, input from orientation sensor(s) **168** may initiate the transition from a two-dimensional map view to a three-dimensional map, and vice versa. For instance, one or more of orientation sensor(s) **168** may detect a tilt (e.g., a user-initiated tilt) in the orientation of the multifunction device and, in response, initiate the aforesaid toggling.

Map **314** may include a graphical position indicator **316**, which may represent the location of the multifunction device within the geographic region of the map. Generally position indicator **316** may represent the current or real-time position of the multifunction device, although it should be understood that in some cases there may exist some small amount of temporal latency between the actual position of the multifunction device and the graphical representation of that location (e.g., position indicator **316**). This may occur, e.g., when the multifunction device is in motion. In various embodiments, the multifunction device may be configured to perform map matching including but not limited to aligning a sequence of observed user positions with a road network on a digital map. In various embodiments, the multifunction device may be configured to perform a “snap to” function in which the graphical position indicator **316** is aligned onto a roadway when the user’s position falls within a specified threshold distance of the roadway.

Furthermore, multifunction device **300** may generally be operated by a user. For example, multifunction device **300** may in some cases be a smartphone utilized by an individual to make phone calls, send text messages, browse the internet, etc. As use of multifunction device by an individual generally implies the individual is proximate to the multifunction device (e.g., the user may be holding the device in his or her hand), references herein to the location of the device and the location of the user may be considered to be synonymous. However, it should be understood that in some cases the actual position of the multifunction device and the user of that device may differ by some distance. For instance, the user may place his or her multifunction device on a table of an outdoor café while sitting in a nearby chair. In this case, the position of the device and the position of the user may differ by some small amount. In another example, multifunction device **300** may be mounted on a car dashboard (e.g., for use as a navigation device) while the user of the device sits nearby (e.g., in the driver seat of the car). In this case as well, the position of the device and the position of the user may differ by some small amount. Despite these small differences in

position, generally the position of the multifunction device and the position of the multifunction device user may be considered to coincide.

In various embodiments, the map **314** displayed by the multifunction device may include one or more roads (e.g., roads **318a-b**), buildings (not illustrated), terrain features (e.g., hills, mountains) (not illustrated), parks (not illustrated), water bodies (not illustrated), and/or any other item that may be conveyed by a map. In some cases, the map may also include other map or navigation information including but limited to readouts from one or more of a directional compass, an altimeter, and/or a thermometer.

In various embodiments, the mapping application may be configured to generate directions from an origination (e.g., an address or a user’s current position) to a destination (e.g., an address, landmark, bookmarked/saved location, or point of interest). For instance, an indication of the origination and/or destination may be input into the multifunction device by the user. The multifunction device may generate one or more candidate routes between those two points. The multifunction device may select one of those routes for display on the device. In other cases, multiple candidate routes may be presented to the user and the user may select a preferred route. In the illustrated embodiment, one route is illustrated as route **320**. The route may also include turn-by-turn directions which may be presented to the user (in 2D or 3D), such as a graphical indication to perform a turn **322a** from road **318a** to road **318b**. In some embodiments, this graphical indication to perform a turn may be supplemented or substituted with an audible indication to turn, such as a voice command from speaker **310** that indicates the user is to “turn left in 100 yards,” for example. In some embodiments, the route that is selected may be presented to the user as a route overview. For instance, before proceeding with navigation, the multifunction device may generate a route overview display that graphically indicates key information for the route, such as key turns, route distance and/or an estimated time for traversing the route. In some cases, the multifunction device may be configured to generate a display of driving maneuvers (e.g., turns, lane changes, etc.) that occur in quick succession, either in the route overview or during actual navigation. This information may help the user safely prepare for such maneuvers. In some cases, the route information may be presented in a list format, such as a list of turns or other maneuvers.

In various embodiments, the mapping application of the multifunction device may be configured to track the position of the user over time and correspondingly adjust the graphical position indicator **316** to indicate the new position. For instance, the mapping application may determine that the user is traveling along route **320** from position information (e.g., information from GPS module **135**) and update the map **314** accordingly. For instance, in some cases the map **314** may remain stationary while position indicator **316** is moved along the route. In other cases, position indicator **316** may remain stationary or “fixed” while map **314** is moved (e.g., panned, turned, etc.) around the position indicator.

In various embodiments, the multifunction device may be configured to display alternate or contingency routes. In some cases, these routes may be selectable by the user (e.g., via the touch screen interface). In other cases, the multifunction device may select a best route based on one or more parameters, such as shortest distance or time. In some cases, these parameters or preferences may be set by the user.

As described in more detail below, the multifunction device may in various embodiments receive routing information that specifies a route from a map service. In some case, the multifunction device may carry out navigation guidance

in accordance with this route. However, in some cases, the multifunction device may perform a reroute operation in order to generate a new route to the destination. For instance, the user may have deviated from the original route or explicitly requested a new route. In some cases, the multifunction device may perform rerouting based on cached map data stored on the multifunction device.

In various embodiments, the multifunction device may be configured to perform route correction based on real-time data, such as updates in map information, road conditions, traffic conditions, and/or weather conditions. For instance, the multifunction device may be configured to alter a route such that the route avoids a construction zone or a dangerous storm cell.

In various embodiments, the multifunction device may be configured to perform lane guidance independently or as part of navigation guidance. For instance, the multifunction device may, in response to detecting that multiple turns follow in quick succession, provide the user with a direction or suggestion as to which lane to occupy. For instance, a voice or visual indication may specify that the user “turn right, then move to the left lane” in anticipation of a subsequent left turn. In another example, the multifunction device may detect one or more lane closures (e.g., due to construction or other reasons) and instruct the user to avoid such lanes.

In various embodiments, the multifunction device may be configured to generate voice prompts for directions. For instance, during navigation guidance, the multifunction device may be configured to generate audio representations of the next turn or driving maneuver on the route. For instance, the multifunction device may be configured to audibly indicate the user should “turn left in 100 yards” or some other audible indication of a maneuver.

In various embodiments, the multifunction device may be responsive to various voice commands for performing actions including a command to obtain a route. For instance, the multifunction device may interpret the user’s voice through a microphone or other transducer of the multifunction device. The user may specify an origination and a destination for the requested route. In various embodiments, the multifunction device may be configured to utilize the user’s current location as the origination for the route.

In various embodiments, the multifunction device may be configured to perform a search along a specific route, such as current navigation route. For instance, the user of the multifunction device may request the location of points of interest, such as fuel stations or restaurants. However, if a user is traveling along a particular route, they may not be particularly interested in points of interest that are not proximate to that route. As such, the multifunction device may be configured to scope any searches to points of interest within a specified distance away from the route. In various embodiments, this distance may be a configurable parameter.

In various embodiments, the multifunction device may be configured to display various graphical layers including but not limited to a graphical map information, aerial images (e.g., satellite-acquired images), and/or traffic information. For instance, in the traffic information example, the multifunction device may overlay color coded traffic information on roadways to indicate the speed at which traffic is flowing. For example, green color coding may be used to indicate traffic is flowing normally, and yellow or red may be used to indicate traffic slowdowns.

In various embodiments, the multifunction device may be configured to display any quantity of metrics or statistics about a navigation route including but not limited to an esti-

mated time of arrival, travel distance remaining, average speed (overall or moving average), top speed, and/or other route statistics.

In various embodiments, the multifunction device may be configured to display routes at different angles in order to accommodate the preferences of different users. Such viewing angles may include a birds eye view for two-dimensional maps to any of a variety of camera angles available for a three-dimensional map.

In various embodiments, the multifunction device may be configured to provide navigation information other than map and routing information. For instance the multifunction device may expose output from any of the hardware device described above with respect to FIG. 1. In one non-limiting example, an orientation sensor 168 may include a compass that outputs direction data. The multifunction device described herein may be configured to display this directional data as a virtual compass, for example.

Map Tool

FIGS. 4A-4F are flowcharts depicting selected processing stages of embodiments of a map tool as implemented within a mapping application. The mapping application, as discussed above, may be invoked through the a user selecting the mapping application through the interface of a multifunction device 300. The mapping application may engage the services of the map service operating system as described in regard to FIG. 9.

In some embodiments, given information on the 3D space of the surrounding environment, the coordinates or location of a given label in the 3D space, and a virtual camera viewpoint in the 3D space, a map tool may determine how occluded the given label is relative to the coordinates of the virtual camera and objects in the 3D space.

In some embodiments, the map view of a map region to be displayed on a mobile device may be represented with a two-dimensional grid corresponding to the map region. For example, if the map region to be displayed is a 100 square meter area, a grid line in the grid representation may be defined at each meter, resulting in 10,000 grid segments each the size of one square meter. Given the 10,000 segment grid, vector-based mapping information such as 3D mesh data may be referenced in order to correlate the 3D mesh data for the map region to the grid representation of the map region. Once the representation of the grid has been matched to corresponding mapping information in the 3D mesh, a height for each grid segment may be determined, where the height information is drawn from the 3D mesh data for a given grid location. In this way, the grid representation may serve to construct a 3D model approximating the 3D objects within the 3D mesh data. In this embodiment, the constructed 3D world is a collection of square meter polygons extruded to various heights extracted from the 3D mesh data. In other cases, the grid area may be divided into more or fewer grid segments in balancing between computational complexity and speed.

Further, depending on the size of the grid segments, there may be more or less accuracy of the 3D model representing the physical world of the map region. In other words, it may be the case where one segment includes a corner of a building, and the segment is determined to be the height of the building, however, part of the segment is not occupied by the building yet the entire segment is assigned to be the height of the building. In this way, a measure of occlusion may determine that part of a label is occluded based on the segment protruding to the height of the building. However, the occlusion is a product of the inaccuracy of the model rather than the label

actually being occluded. As noted above, the degree of granularity may be selected based on considerations to balance accuracy with speed.

Given the coordinates of a label location within the map view, coordinates for the label may be determined with respect to the 3D model, and one or more measures of occlusion may be performed to determine whether or not the label should be rendered in a map view. In different embodiments, a grid representation may be subdivided into different sizes or granularities. Further, the grid shapes may be different geometric figures, such as triangles, hexagons, or ovals. In some embodiments, the constructed 3D model does not provide information regarding what type of object exists at each grid section. Further in this embodiment, the 3D model provides rough information regarding the boundaries of any objects or structures in the 3D model. As noted above, a building may end at the beginning of a grid segment, however, in the model, the building takes up the entire grid segment. However, as discussed below, the measures of occlusion take into account any inaccuracies in the boundaries of objects in the 3D model.

As per FIG. 4C, in some embodiments, the map tool may generate a representation of a map region, where the representation is divided into multiple sections, as described above, and as reflected in stage 442. The map region corresponds to the area of a map that serves as the basis of a map view.

The representation of the map region is initially only two-dimensional, representing the surface of the map region. To create a 3D model, the map tool may extrude, or add a height dimension to each of the sections of the representation of the map region. In order to add the height dimensions to the representation of the map region, the map tool uses height information received that corresponds to the map region, as reflected in stage 444.

Given the height information, the map tool may then extrude each respective section of the representation of the map region according to a corresponding height value of the height values in the received height information, as reflected in stage 446. Further in stage 446, the map tool uses a correspondence between a respective section of the representation of the map region to a respective location in the map region, where the respective location in the map region corresponds to a respective height value. For example, say that a section in the representation of the map region corresponds to a location in the map region with certain coordinates, in this case, the height values may be organized such that a given location coordinate corresponds to a particular height for an object at that location coordinate. In this way, a particular height value for a given section of the representation of the map region may be determined from the multiple height values.

At this point, the map tool has both 2D information in the form of the representation of the map region and 3D information in the form of height values. Using the 2D and 3D information the map tool may generate a 3D representation of the map region based on extruding each respective section of the representation of the map region according to a respective height value of the height values, as reflected in stage 448.

Based on the 3D representation of the map region, the map tool may determine that a particular label that is to be rendered in a map view is occluded below a particular threshold amount, as reflected in stage 450. A measure of occlusion for the label may be determined through any of the previously described methods, and the threshold may be adjustable to allow for variability in the amount of occlusion to be tolerated.

Based on the determination that the label in the map view of the map region is occluded below the particular threshold,

the map tool may then render at least a portion of the label in the map view, as reflected in stage 452. Once the map tool similarly processes each label within the map region, the map view rendered in the display of multifunction device 300 would only include labels that are not occluded beyond the tolerance indicated by the threshold value.

In some embodiments, upon the mapping application being invoked, the map tool may access, receive, or request mapping information for the current location of the multifunction device 300. As discussed above, in some cases, the mapping application may periodically or aperiodically receive GPS information for the current location of the multifunction device. In other cases, in response to the mapping application being invoked, the map tool may request or receive 2D mapping information and also receive 3D mesh data from a server such as Map Service 930, or from some other source, where the 2D mapping information includes footprint information for buildings within or partially within the map region.

In some embodiments, the map tool may construct a 3D model representative of the surrounding environment, and specifically, representative of the map region to be displayed in the map view. The 3D model may be constructed based on an approximation of the 2D and 3D mapping information for the map region.

In some cases, the 3D model may be constructed through the use of a grid representation of the map region and the mapping information. For example, for the current map region, the map tool may partition the representation of the map region into a grid, where the area of the grid may be overlaid, or otherwise mapped or correlated to a respective area of the map region within the 2D mapping information. In this way, footprints of structures from the 2D mapping information may be correlated to one or more grid segments.

Further, footprints of other objects in the map may similarly be correlated to one or more grid segments, for example, trees or works of art. In this example, to generate a 3D model from the 2D information in the grid, the map tool may access the 3D mapping information. For example, for a given footprint in the grid, the 3D mapping information may be accessed to determine one or more points corresponding to the footprint and the map tool may further determine the heights of the one or more points corresponding to the footprint. In this example, multiple height values for multiple locations in the footprints may be correlated to the 3D information, and the highest height value may be selected.

In some embodiments, a footprint may be divided into multiple regions, and height values may be correlated to the 3D information for one or more locations within each of the regions. For example, a building may be L-shaped, where one leg of the L is not as tall as the other leg of the L. In this case, the footprint may be divided into two regions, one region corresponding to each leg of the L, and a height value for each region may be used as the basis for extruding the footprint into three dimensional space. The regions of the footprint may be determined in other manners, such as uniform division of the footprint into any given number of regions, for example, dividing each footprint into quarters.

In some cases, given the height information for the one or more points for each footprint in the grid, the map tool may extrude the footprint into the third dimension based on the height information. In some cases, the height of the extruded footprint may be based on the highest of the points corresponding to a particular footprint. In other cases, the height of the extruded footprint may be based on an average of the majority of the points corresponding to the particular footprint. For example, if five points of height data corresponding

31

to a particular footprint are determined from the 3D mapping information, then the average height of three of the highest points of height data may be determined to be the height of the extruded footprint. In other cases, the single highest point may serve as the basis for the extruded footprint. In some cases where the footprint is not a polygon, the footprint may be modified to a polygon.

Other calculations and embodiments are possible where a height may be determined for a particular footprint based on 3D mapping information, and this general concept of approximating one or more spatial dimensions for one or more objects in the map, where the approximations are based on mapping information.

In some cases, the spatial dimensions may be based on only the 2D information. For example, the map tool may base the 3D model on the performance capabilities of the multifunction device 300. For example, if multifunction device 300 is contains an older model processor and limited memory resources, then the map tool may only use the 2D mapping information and extrapolate each 2D footprint in the map to infinity, or some height high enough to obstruct, or the maximum height. In other cases, without extrapolating a 2D footprint, the map tool may simply determine, as part of a measure of occlusion, whether a label is beyond any of the points within the footprint for a given object, and if so, whether a given point on a label from which a ray is projected in 2D space intersects with any of the points in the given 2D footprint. While the amount of error increases using only 2D information, the tradeoff in speed may be preferable to waiting for a label to be rendered. In the case that the map tool determines that a multifunction device has a more capable processor, or hardware resources, or available crowdsourcing computation, then the map tool may perform a more computationally intensive measure of occlusion, using a more detailed 3D model of the surrounding environment.

In this example, given the constructed 3D model, or more generally, the one or more spatial dimensions for one or more objects in a map, the map tool may determine a measure of occlusion for a label in the map, where the determination is based on the above approximations.

The mapping information used within the map tool may include labels for various elements within the map area, such as points of interest, street names, landmark names, or any other type of annotation to the map. As illustrated in FIGS. 5A and 5B, the map tool of the mapping application may render or cause to display the map views as depicted within multifunction device 300 screen 312.

As per FIG. 4A, in one embodiment, upon the mapping application being invoked, the map tool may access, receive, or request mapping information for a map region of the current location of the multifunction device 300, as depicted in stage 402. The mapping information may include height information for various objects within the map region. For example, the map tool may receive three-dimensional mapping information for the map region from map service 930 or from GPS 940. For example, in some cases, the mapping application may periodically or aperiodically receive GPS information from GPS 940 and/or 2D or 3D information from Map Service 930. In other cases, in response to the mapping application being invoked, the map tool may request 3D mesh data or 2D data from a server, or from some other source. In some cases, the mapping information is vector data and not raster image data.

In some embodiments, the multifunction device may also receive crowdsourced information. For example, in some cases, multifunction device 300 may provide the user with the ability to enter information regarding the map information.

32

For example, in a given map view, for example, map view 312 of FIG. 5A, a user may tap and hold over the region proximate to label 514, or proximate to the object. In response to the tap and hold over label 514, the map tool may provide the user with the option to indicate that the labeled object no longer exists. A server in communication with multifunction device 300 may receive the information and may update the mapping information corresponding to the object. In some cases, to minimize the possibility of error of relying on a single source, the server may update the mapping information after some threshold of users have indicated the same update to a given object. The updated information may then be provided to other mobile device users receiving mapping information.

In some embodiments, if the map tool is basing a measure of occlusion on 3D information that is highly detailed or complex, the map tool may partition the problem of calculating portions of the measure of occlusion to other devices in range of the wireless capability of the multifunction device. For example, if a user is waiting for a light standing next to another user on a street corner, or walking alongside the other user, and the other user has enabled their device to openly communicate wirelessly with other devices, then the multifunction device for the user may assign a portion of the calculation of the measure of occlusion to the other device. For example, a calculation of whether one or more projected rays from the virtual camera viewpoint collide with any objects before arriving at a given label. Similarly, the multifunction device may communicate with multiple other devices. In other cases, the multifunction device of the other user may have already performed a calculation of a measure of occlusion, and when the multifunction device arrives to the same spot, or very near the same spot of the other user, the other multifunction device may simply transmit an already performed measure of occlusion—thereby allowing the newly arrived user of the multifunction device to avoid performing a determination of a measure of occlusion. In this way, the multifunction device may make use of crowdsourcing information from random collection of other users that happen to be in the range of the multifunction device. In some cases, the wireless communication may occur via Bluetooth™

In other embodiments, a user may delete objects in the map view, and once deleted, the map tool may determine whether previously occluded label is now visible, or more visible. For example, if a user looking at map view 312 of FIG. 5B sees that street 522 intersects street 502, and that street 502 is not labeled, the user may delete one or more of the obstructions. In some cases, a user may tap and hold over the region of the map view corresponding to an object, and in response, the map tool may provide the user with the option to remove the object from view. Once the object is removed, the 3D model is updated, and the updated 3D model may provide a different measure of occlusion than in the 3D model with the object. In this example, given the updated measure of occlusion for a label corresponding to street 502, the label corresponding to street 502 may then be rendered, or drawn within the map view. In other cases, the label may be determined to fit within the area display area emptied through the removal of an object from the map view, and therefore the label may be moved to the open area before performing a determination of occlusion.

In some cases, based on a projection of where the user is going, the map tool may prefetch corresponding mapping information, and if the projection is wrong or partially wrong, some or all of the prefetched mapping information may be discarded. In some cases, in addition to the 3D data obtained, the map tool may retrieve from locally stored information, or request from a remote system, 2D mapping information. In

some cases, if a user has entered a destination address, the map tool may determine which mapping information may be necessary after the current mapping information is received, and may request the mapping information to begin generating a map view that is to be displayed at a future time.

Given received mapping information including height values, along with any additional information, the map tool may generate a 3D model of a map region of the surrounding environment, as reflected in stage 404. For example, the height values may be used to determine the volume of a section in a grid representation of the map region. In this example, the map tool may determine whether the volume of any sections in the 3D model are between the line of sight of the virtual camera viewpoint and the label in the 3D model. For example, the 3D model may include a given section from among multiple sections of the grid representation of the map region, where the given section may have boundaries and a volume such that a ray projected from a virtual camera viewpoint for a map view would collide with the given section. This determination may occur for each label that is to be rendered in the map view of the map region.

Based on the 3D model of the map region and based on coordinates for a label to be displayed in a map view of the map region, the map tool may determine a measure of occlusion for the label as it would be displayed in the map view, as reflected in stage 406. Further as reflected in stage 406, the coordinates for the label as it would be displayed in the map view correspond to coordinates for the label in the 3D model.

The measure of occlusion may include various calculations or determinations regarding the visibility of the label with respect to the virtual camera viewpoint of multifunction device 300. The measure of occlusion may provide the map tool with information to determine whether or not to render a given label or at least a portion of the given label within the map view. In response to the measure of occlusion for a label satisfying a threshold to be rendered in the map view, the map tool may then render the label or a portion of the label in map view of the map region, as reflected in stage 408.

In some embodiments, a user may manipulate a given map view such as through input indicating to the map tool to display a different virtual camera viewpoint of the map view. For example, within a given map view, a user may wish to see the map view from the other side of a building. As a result, the amount of occlusion for one or more labels within the map view may change and new calculations of occlusion may be performed to determine whether or not a label is occluded from the new virtual camera viewpoint. However, given that the previously generated 3D model of the map view has already been constructed, the measure of occlusion does not need to generate a new 3D model of the map view because the locations and spatial dimensions of object in the previously generated 3D model remain valid for the new virtual camera viewpoint.

As per FIG. 4B, in some embodiments, upon the mapping application being invoked on a multifunction device 300, the map tool may access, receive, or request height values for a map region or access, receive, or request mapping information for the map region that may include height values for a current location and virtual camera viewpoint for the multifunction device 300. In this example, the reception of mapping information including height values for a map region.

Based on the mapping information, the map tool may construct a three-dimensional model approximating the map region, as reflected in stage 424. The map tool may construct the three-dimensional model in any number of ways, including each of the methods discussed above. For example, the map tool may use a grid representation of the map region and

extruding sections of the grid representation based on the height values for the map region in order to create a 3D version of the grid representation.

The map tool may determine a measure of occlusion for a label to be displayed in a map view of the map region, where the determination is based on the 3D model of the map region and also based on a correlation between the 3D model of the map region and the map view of the map region, as reflected in stage 426. The determination of the measure of occlusion may be performed in any of the embodiments discussed in other sections.

Based on the measure of occlusion for the label in the map view of the map region being less than a threshold value, the map tool may render the label or at least a portion of the label in the map view, as reflected in stage 428. The map tool may repeat this process for each label that is to be displayed in the map view. In this way, for a given map view that would otherwise include rendered labels with difficult to discern information, the map view is rendered such that only labels with useful information, or labels that are not too occluded, are rendered within a map view seen on the multifunction device 300.

FIG. 5A depicts a map view as seen from the virtual camera viewpoint of the multifunction device 300, which the map tool uses as a proxy for the user's viewpoint at ground level.

FIG. 5B depicts a map view for a virtual camera viewpoint as seen from a bird's eye viewpoint that corresponds to a point above and slightly behind the multifunction device 300. In both cases, the determinations of whether or not a label is occluded is made from the respective point of view.

In FIG. 5B, tree 504 and buildings 506 and 508 are obscuring street 502, where street 502 corresponds to a label "Martin Luther King, Jr. Way". Further, within the map view, labels 510, 512, 514, and 530 are displayed corresponding to, respectively, street 522, street 524, structure 508, and structure 506. However, if a label for a street, street 502, were rendered within the map view, in this map view most of the label may be occluded, and only a portion of the label may be displayable. For example, street 502, if rendered within the map view, would only have a small amount of space within which a label may be rendered, the area corresponding to area 526 and the area corresponding to area 528. In this example, neither area 526 nor area 528 is large enough for the entire label, and therefore, only a portion of the label may be displayed. In some cases the portion displayed is not enough to provide meaningful information, for example, street 502 has label "Martin Luther King, Jr. Way" and if only "ay" were displayed, or some other, equally indecipherable portion were displayed, the user would be presented with useless information. On the other hand, in this example, if the entire label were displayed, the label may be displayed over some other object or near enough to some other object such that confusion may arise as to which label corresponds to which object. In such a case, it may be preferable to not display any portion of the label at all.

In some embodiments, the map tool determines a measure of occlusion for a label in the map region, where the determination is made based on a 3D model of the map. In some cases, the 3D mapping information used to determine the measure for occlusion for a label may be based on high-resolution 3D mapping information, such as 3D mesh data. However, most mobile devices have battery or computational limitations that do not allow for responsive, or real-time, displays of map views possible if high-resolution 3D mapping information is used. In any case, so long as the map tool

has access to or generates spatial information for the map, a determination of a measure of occlusion may be performed, as reflected in stage 456.

In some embodiments, the measure of occlusion is based on calculations of a ray projected from a virtual camera viewpoint the 3D model space from the multifunction device 300 to one or more points on the label in the 3D model. In this calculation, the label may have corresponding spatial coordinates specifying where in the 3D model of the map view the label is to be displayed. In an analogy to the real world, the label may be considered the street name on top of a pole located at one of the corners of the intersection of streets. In this analogy, if a user is standing 51 meters away from the intersection, one or more objects may obstruct the view of the user, for example, the corner of a building may be the obstruction. In a similar way, if the label is located in the virtual 3D model as the street sign is located in real space, the virtual building may obstruct the view of the label as it would be displayed in the map view according to the spatial coordinates for the label.

In such a case, the above mentioned ray that is projected out from a virtual camera viewpoint corresponding to the multifunction device within the 3D model is targeted to a point in the label. If the ray reaches the label without colliding or intersecting with any object in the 3D model, the ray is considered to have successfully reached the label. The result of the ray "hitting" the label results in a data point that may be used in calculating an overall measure of occlusion. In some embodiments, multiple points in the label are targeted by the projected ray, and each projected ray corresponding to a data value indicating whether or not the ray reached the label without going through the space defined by an object in the 3D model.

In some embodiments, the calculation of the measure of occlusion may be generated as a measure of the number of projected rays that reached the label without intersecting an object in the 3D model. For example, eight rays may be projected, one ray for each corner of the label and two rays along the top perimeter of the label placed at the $\frac{1}{3}$ and $\frac{2}{3}$ points along the length of the top, and two rays along the bottom perimeter of the label also placed at the $\frac{1}{3}$ and $\frac{2}{3}$ points along the bottom. In this example, if the two leftmost rays are blocked by objects in the 3D model and the remaining six rays successfully reach the label without obstruction, then the measure of occlusion may produce a low value, corresponding to the small amount of occluded area of the label. This calculation may equivalently be framed in terms of a measure of visibility, in which case the calculation would be the same and the measure of visibility would be the inverse of the above-mentioned measure of occlusion.

In other embodiments, more or fewer rays may be projected, and different heuristics may be applied to interpret the results of the projected rays. For example, three rays may be projected along a center line dividing the label into a top half and a bottom half, with target points along the two edges and the center, and in this example, if a projected ray hits any two target points, the measure of occlusion may be determined to be a small value, resulting in the label being drawn.

In some embodiments, for any given ray projected from the multifunction device to the label, the measure of occlusion may factor in how many, and/or which types of objects were hit by the projected ray. For example, instead of simply counting up how many projected rays reached a target point on a label, each projected ray may be weighted. For example, if a projected ray intersects with a tree, the tree may be assigned a weight corresponding to a low obstruction value because trees sometimes allow some visibility to objects on the other

side. By contrast, a building may be assigned a weight corresponding to a high obstruction value because most buildings do not allow visibility to objects on the other side. In this way, if a projected ray to a target point on a corner intersects with a tree, and a projected ray to a target point on another corner also intersects with a tree, and a projected ray to a middle target point on the label intersects with nothing, the measure of occlusion may be low based on the weighted assignment of low values to the tree intersections. However, in this same example, where trees are replaced with buildings, the measure of occlusion may be significantly higher based on weighted assignments of high values to the projected ray intersections with buildings.

In some embodiments, where the map tool is operating based on a low resolution model of the surrounding environment, the accuracy of the exact boundaries of a given building or object in the 3D model may be low. In other words, if a projected ray from the multifunction device to a point on a label intersects with only one object or structure on the way to the target point on the label, the intersection may be weighted to minimize the value of the single collision. However, in this example, if the same projected ray intersected with two objects in the 3D model, the weight corresponding to intersection calculations may be a high value, indicating high confidence in an actual obstruction. In other cases, instead of weights, if a projected ray intersects with two or more objects, the projected ray may be determined to not have reached the target point at all.

In some embodiments, the measure of occlusion may be based on the granularity of the 3D model. For example, if the granularity of the 3D model into which a ray is projected has a high degree of detail, then the measure of occlusion may be based on a single object collision of the ray projected toward the label instead of multiple object collisions. For example, using 3D mesh data directly, or constructing a 3D model to include additional details regarding dimensions of objects in the 3D model. In other cases, if the granularity of the 3D model into which a ray is projected has a low degree of detail, then the measure of occlusion may be based on multiple object collisions of the ray projected toward the label. In other words, with low granularity 3D models, there is lower confidence that a calculated collision between an object in the 3D model and a projected ray actually corresponds to a projected ray colliding with that same object in the real world. In some cases, the granularity of the 3D model used may be based on the available computational cycles available to the map tool. In other words, if the multifunction device has few computational and/or resource demands from other applications, then more computational cycles may be available to the map tool.

Further in the cases of using different resolutions of 3D models, calculated collisions may be given greater or less confidence based on which part of an object has been struck with a projected ray. For example, in a low resolution 3D model, if a projected ray collides with the outer edge of an object, the collision may be given a small amount of weight based on the lower degree of confidence that the edge of a 3D model object is where the edge of the corresponding actual object is located in the real world. Similarly, in a high resolution 3D model, if a projected ray collides with the outer edge of an object, the collision may be given a large amount of weight based on the high degree of confidence that the edge in the 3D model object is where the edge of the corresponding actual object is located in the real world.

In some embodiments, another factor in the determination of a measure of occlusion is the density of the surrounding environment. For example, within an urban environment, if there are many tall objects in the map, such as tall structures

on each city block, then a density value may be set to indicate a high level of density. Similarly, the density value may reflect a low level of surrounding density of objects. Given a measure of density, the measure of occlusion may be adjusted toward indicating not to render a label if density is high, or if the density is beyond a threshold value. Similarly, a density value may be used as a factor in assigning the weighted values discussed above.

In some embodiments, a user may prefer to see partially occluded labels displayed even if the result is additional clutter in the map view. Similarly, in some cases, the user may prefer to see fewer partially occluded labels displayed in order for the map view to appear cleaner, or less cluttered. In either case, one or more of the factors on which a measure of occlusion is based may be adjusted by a user to achieve either displaying more or fewer partially obscured labels.

In some cases, the user may simply adjust a top-level user setting corresponding to low, medium, and high levels of acceptance for partially displayable labels. In other cases, a user may be presented with an option to display a lower or higher resolution 3D model in the map view. In this case, the 3D model displayed may be the 3D model generated, and given the higher granularity, the map tool may adjust the weights for the confidence values in calculating collisions of objects and rays. Further in response to a setting of accepting a low, medium, or high level of occlusion for labels in the map view, the map tool may correspondingly adjust whether a greater or fewer number of calculated collisions may be acceptable before determining whether or not a label in the map view is considered occluded enough to not be displayed. Other factors in determining a measure of occlusion may be similarly modified in response to the user settings.

In some embodiments, the measure of occlusion may, in addition to calculated values discussed above, may also be based on other factors related to the display of a map view in display area 312 of the multifunction device. For example, if the distance between the projection point (the user/multifunction device position) and the label is large, there may be less urgency in displaying a label, and therefore, if the measure of occlusion is not strongly indicative of displaying a label, then based on the large distance, the measure of occlusion may be overridden to not display the label. In other cases, if the distance between a label is small or rapidly shrinking, then a measure of occlusion that weakly indicates not displaying a label may be overridden to display the label—on the basis that the label has greater importance because it may be useful for a user to make a navigation decision in the near future.

Given the calculated measure of occlusion and other potential considerations, the map tool renders the label in the map, where the decision to render the label is based on the measure of occlusion, and where the decision to render the label may further be based on other potential considerations, as depicted in stage 458. In some cases, the decision to render the label may be in response to the measure of occlusion being greater than some threshold value, where the threshold value may be based on a default setting defined prior to the invocation of the mapping application, or based on a user-defined setting.

In some embodiments, the measure of occlusion may reuse data from previous determinations of occlusion or from previous calculations. For example, as a user walks down a street, new mapping information may be incrementally introduced into the navigation area. Given the overlap of previous mapping information to newer mapping information, the map tool may reuse, for example, determinations of spatial dimensions for objects that remain in the mapping information from a previous calculation based on prior mapping information to a current determination of occlusion based on new mapping

information. Given that the spatial dimensions would not change, it is an increase in calculation efficiency for the map tool to reuse previously generated information.

In some embodiments, the map tool may further base a measure of occlusion on the amount of zoom into a map view. In other words, in this embodiment, the map tool may base a measure of occlusion on any combination, or on any subset of: (a) density of the surrounding environment, (b) virtual camera viewpoint, or (c) zoom levels. For example, if a user zooms into a map view a given displayed object may take up more of the display area 312, and consequently, a given object may be displayed with more pixels. In such a case, the granularity of the visible is finer, and the measure of occlusion may be adjusted as discussed above with regard to different granularities of a given 3D model. In this example, if a label is displayed at the same size, irrespective of zoom level, then upon a zoom a given label may become larger relative to a given object. In this case, the calculations of projected rays may change because one or more targeted points may now be reachable with a projected ray without any collision with a given object.

In one embodiment, the map tool may receive an indication of a map region. In response to the indication of a map region, the map tool may request mapping information corresponding to the map region. The navigation may then generate a three-dimensional model based on the mapping information, where the 3D model includes multiple objects. Once the 3D model is constructed, the map tool may proceed to identify a given label from among the labels in the 3D model, where the identified label has coordinates placing the label in a particular location in the 3D model. With respect to a viewpoint from which the map view is generated, one or more points in the label are determined to be within an unobstructed line of sight from the viewpoint. Given a determination of how many lines of sight to the various points on the label are unobstructed, a further determination may be made as to whether enough lines of sight to the label are within a clear line of sight so that rendering the label provides useful information to a user. In this case, at least a portion of the label is rendered, where the decision to render the label is based on the determination of the one or more points in the label that are within a clear line of sight from the viewpoint. Less than the entire label may be rendered if enough lines of sight were clear such that the visible portion of the label would be useful, but at least some parts of the label are obstructed.

As per FIG. 4D, in some embodiments, to generate a rendering of a map view of a map region, the map tool may receive a set of mapping information for a given map region, as reflected in stage 462. The map tool may also receive another, different set of mapping information for the given map region, as reflected in stage 464. The map tool may then correlate the first set of mapping information with the other, second set of mapping information, as reflected in stage 466. The correlation of data may allow for the generation of a three-dimensional model with only a selected level of granularity, where the level of granularity may be different than the granularities of the two sets of mapping information received. For example, if one set of mapping information is 2D information and the other set of mapping information is highly detailed 3D information, a three-dimensional model may be constructed based on correlating common elements from the two data sets and based on the virtual camera viewpoint from which the map view is to be generated, as reflected in stage 468.

In some cases, additional data sets may be received and correlated to generate the 3D model. Once the 3D model is constructed, the map tool, based on a virtual camera view-

point, may proceed to calculate one or more measure of occlusion for a given label in the 3D model space, as described above in different embodiments and as reflected in stage 470, where the calculation is further based on a correlation between the 3D model and the map view of the map region. Given the one or more measures of occlusion, the map tool may then determine to render a given label in the map view of the map region, as reflected in stage 472.

As per FIG. 4E, in some embodiments, the map tool may receive mapping information for a map region for which a map view is to be rendered, as reflected in stage 484. The mapping information may be any combination of mapping information described in other embodiments, and the map tool may then generate a 3D model of the map region based on the received mapping information for the region, as reflected in stage 486.

Given the constructed 3D model of the map region, the map tool may proceed to identify a label that is to be rendered in a map view for the map region, as reflected in stage 488. The label may be identified from the received mapping information which may include information regarding spatial dimensions and coordinates for elements in the map region, and which may also include metadata for some or all of the elements in the map region, such as the coordinates of labels and information on the names of structures or parks, or information on aggregated user rankings for a restaurant within the map region.

From the identified label, the map tool may use the coordinates of the label within the map region in order to determine corresponding coordinates for where the label would be within the 3D model, as reflected in stage 390. In the cases where the area of the map region for the 3D model is the same area of the map region of the mapping information, the coordinates for the label would be the same.

The map tool may then determine that one or more points within the space of the label coordinates within the 3D model are unobstructed, as determined from a respective ray projected from a virtual camera viewpoint to a respective point within the space of the label coordinates reaching the respective point without intersecting or colliding with any object in the 3D model space. The determination of whether one or more points in the label coordinates are unobstructed is reflected in stage 492.

As discussed in other embodiments, the determination of whether or not to render a label may depend on how much of the label is occluded, and how much of the label is occluded may be based on how many points within the label are unobstructed. In some cases, a threshold amount of points in a label may be defined such that if the threshold amount of points in the label are unobstructed, the label is determined to be fit for rendering. In this case, as reflected in stage 494, the map tool renders at least a portion of the label in the map view based on the determination that the one or more points in the label coordinates satisfy the threshold amount of points.

As per FIG. 4F, in some embodiments, the map tool may receive two-dimensional mapping information for a map corresponding to a map view to be displayed in a mapping application, as reflected in stage 495. The navigation tool may then determine coordinates for a virtual camera viewpoint within the two-dimensional mapping information, corresponding to the user's relative location within the map, as reflected in stage 496. The navigation may then calculate, based on the virtual camera viewpoint, one or more measures of visibility, or measures of occlusion, of the label within the two-dimensional mapping information, as reflected in stage 497.

In this example, using only two-dimensional mapping information, a decision on whether a given point in a label is visible from the virtual camera viewpoint may be less nuanced. In other words, in the two-dimensional case, if a line of sight, or projected ray, from the virtual camera viewpoint to a point on a label intersects with a two-dimensional footprint representative of any object in the two-dimensional map may be considered an obstructed line of sight. In this example, certain objects may be assumed to have sufficient height so as to obstruct a view of a given label, for example, every building may be assumed to have sufficient height, as well as every tree or landmark. However, some objects in the two-dimensional map may be assumed to not be obstructive of a line of sight, for example, benches or utility boxes or flower beds. In this way, the objects in a two-dimensional map are defined to be tall enough to obstruct or not tall enough to obstruct.

As reflected in stage 498, based on the calculation in the two-dimensional space, the map tool determines whether or not to render a given label in the map view of the map. Further in this example, the map view need not be two-dimensional. In other words, the map tool may receive and render three-dimensional mapping information, however, the map tool may simply base all calculations of occlusion on the two-dimensional mapping information. The three-dimensional mapping information may also be referenced to locate a given label in the two-dimensional mapping information since the two-dimensional mapping information, depending on the source, may not include label information. Therefore, in this embodiment, once a label is determined to be obstructed enough not to render, the corresponding label is identified in the three-dimensional mapping information so that the particular label is not rendered.

Example Embodiment

Map Tool Module

FIG. 6 illustrates an embodiment of a Map Tool Module 600. As noted above, the Map Tool Module 600 may be implemented in a variety of different embodiments.

In some embodiments, Control Module 604 may receive Input 602, which may be various types of mapping information. Given the mapping information, Control Module 604 may invoke Model Generation Module 606 to generate a model of the surrounding environment, according to various embodiments discussed above.

Given a model of the surrounding environment, Control Module 604 may invoke Occlusion Measuring Module 608 to determine whether or not a label should be rendered, according to the various embodiments discussed above. This process may be repeated for each label in the surrounding environment within the map.

Given a model of the surrounding environment, Control Module 604 may invoke User Interface Module 608 in response to various user inputs indicating, among other things, a selection of an object within a map view, information for labeling an object in the map view, or requesting more information or specific information regarding an object in the map view. In some cases, depending on the input, Map Tool Module 600 may communicate with Map Service 930 through Map Information Interface Module 610 to request or receive mapping information.

Once Control Module 604 has a determination of whether or not to label a given label in a map, a map view may be rendered reflecting the determination. The rendered map view may be Output 620.

Example Embodiment

Client/Server Architecture

FIG. 7 depicts an illustration, according to one embodiment, of a computing environment in which various devices implementing a map tool may operate. The case of a portable multifunction device has been discussed above.

As another example, a user at a desktop machine may enter any particular address of any given structure, or the user may enter a coordinate, or the user may enter a landmark, or the user may simply navigate to a given map location. Once the user arrives at a map location, the user may navigate as a user on a mobile device would navigate. In this way, without moving, a user at a stationary machine, may navigate through a given map view such that the above-mentioned embodiments of a map tool may operate to present the user with a similar map view as a mobile user would see.

Example Embodiment

Cloud Computing Environment

FIG. 8 depicts one possible computing environment that includes a device **810** accessing a cloud computing environment on a computer **804** over network **802**. In this example, an installed application on a mobile device may be used to access any of the above-discussed embodiments of a map tool executing within one of the virtual computing instances **806** through **808**.

In other embodiments, as discussed above in the case of a client/server architecture, the map tool may be implemented on a computer **812** accessing a cloud computing environment **804** over network **802**. In this example, a user may log in to the cloud computing environment **804** to access a virtual computing instance within which an embodiment of the map tool is executing or may be executed.

Example Computer System

FIG. 10 illustrates computer system **9900** that may execute the embodiments discussed above. In different embodiments, the computer system may be any of various types of devices, including, but not limited to, a personal computer system, desktop computer, laptop, notebook, or netbook computer, mainframe computer system, handheld computer, workstation, network computer, a camera, a set top box, a mobile device, a consumer device, video game console, handheld video game device, application server, storage device, a television, a video recording device, a peripheral device such as a switch, modem, router, or in general any type of computing or electronic device.

In one embodiment, computer system **9900** includes one or more processors **9910a-9910n** coupled to system memory **9920** via input/output (I/O) interface **9930**. The computer system further includes network interface **9940** coupled to I/O interface **9930**, and one or more input/output devices **9950**, such as cursor control device **9960**, keyboard **9970**, and one or more displays **9980**. In some embodiments, it is contemplated that embodiments may be implemented using a single instance of a computer system, while in other embodiments may be implemented on multiple such systems, or multiple nodes making up a computer system, may be configured to host different portions or instances of embodiments. For example, in one embodiment some elements may be implemented via one or more nodes of the computer system that are distinct from those nodes implementing other elements.

In various embodiments, the computer system may be a uniprocessor system including one processor, or a multiprocessor system including several processors (e.g., two, four, eight, or another suitable number). The processors may be any suitable processor capable of executing instructions. For example, in various embodiments, the processors may be general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x86, PowerPC, SPARC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of processors may commonly, but not necessarily, implement the same ISA.

In some embodiments, at least one processor may be a graphics processing unit. A graphics processing unit or GPU may be considered a dedicated graphics-rendering device for a personal computer, workstation, game console or other computing or electronic device. Modern GPUs may be very efficient at manipulating and displaying computer graphics, and their highly parallel structure may make them more effective than typical CPUs for a range of complex graphical algorithms. For example, a graphics processor may implement a number of graphics primitive operations in a way that makes executing them much faster than drawing directly to the screen with a host central processing unit (CPU). In various embodiments, the content object processing methods disclosed herein may, at least in part, be implemented with program instructions configured for execution on one of, or parallel execution on two or more of, such GPUs. The GPU(s) may implement one or more application programmer interfaces (APIs) that permit programmers to invoke the functionality of the GPU(s). Suitable GPUs may be commercially available from vendors such as NVIDIA Corporation, ATI Technologies (AMD), and others.

System memory within the computer system may be configured to store program instructions and/or data accessible from a processor. In various embodiments, the system memory may be implemented using any suitable memory technology, such as static random access memory (SRAM), synchronous dynamic RAM (SDRAM), nonvolatile/Flash-type memory, or any other type of memory. In the illustrated embodiment, program instructions and data may implement desired functions, such as those described above for the various embodiments are shown stored within system memory **9920** as program instructions **9925** and data storage **9935**, respectively. In other embodiments, program instructions and/or data may be received, sent or stored upon different types of computer-accessible media or on similar media separate from system memory or the computer system. Generally, a computer-accessible medium may include storage media or memory media such as magnetic or optical media, e.g., disk or CD/DVD-ROM coupled to the computer system via the I/O interface. Program instructions and data stored via a computer-accessible medium may be transmitted from transmission media or signals such as electrical, electromagnetic, or digital signals, which may be conveyed via a communication medium such as a network and/or a wireless link, such as may be implemented via the network interface.

In one embodiment, the I/O interface may be configured to coordinate I/O traffic between the processor, the system memory, and any peripheral devices in the device, including a network interface or other peripheral interfaces, such as input/output devices. In some embodiments, the I/O interface may perform any necessary protocol, timing or other data transformations to convert data signals from one component into a format suitable for another component to use. In some embodiments, the I/O interface may include support for devices attached through various types of peripheral buses, such as a variant of the Peripheral Component Interconnect

(PCI) bus standard or the Universal Serial Bus (USB) standard, for example. In some embodiments, the function of the I/O interface may be split into two or more separate components, such as a north bridge and a south bridge, for example. In addition, in some embodiments some or all of the functionality of the I/O interface, such as an interface to system memory, may be incorporated directly into the processor.

The network interface of the computer system may be configured to allow data to be exchanged between the computer system and other devices attached to a network, such as other computer systems, or between nodes of the computer system. In various embodiments, the network interface may support communication via wired or wireless general data networks, such as any suitable type of Ethernet network, for example; via telecommunications/telephony networks such as analog voice networks or digital fiber communications networks; via storage area networks such as Fibre Channel SANs, or via any other suitable type of network and/or protocol.

The I/O devices may, in some embodiments, include one or more display terminals, keyboards, keypads, touchpads, scanning devices, voice or optical recognition devices, or any other devices suitable for entering or retrieving data from one or more computer systems. Multiple I/O devices may be present in the computer system or may be distributed on various nodes of the computer system. In some embodiments, similar I/O devices may be separate from the computer system and may interact with one or more nodes of the computer system through a wired or wireless connection, such as over the network interface.

The memory within the computer system may include program instructions configured to implement each of the embodiments described herein. In one embodiment, the program instructions may include software elements of embodiments of the modules discussed earlier. The data storage within the computer system may include data that may be used in other embodiments. In these other embodiments, other or different software elements and data may be included.

Those skilled in the art will appreciate that the computer system is merely illustrative and is not intended to limit the scope of the embodiments described herein. In particular, the computer system and devices may include any combination of hardware or software that can perform the indicated functions, including a computer, personal computer system, desktop computer, laptop, notebook, or netbook computer, mainframe computer system, handheld computer, workstation, network computer, a camera, a set top box, a mobile device, network device, internet appliance, PDA, wireless phones, pagers, a consumer device, video game console, handheld video game device, application server, storage device, a peripheral device such as a switch, modem, router, or in general any type of computing or electronic device. The computer system may also be connected to other devices that are not illustrated, or instead may operate as a stand-alone system. In addition, the functionality depicted within the illustrated components may in some embodiments be combined in fewer components or distributed in additional components. Similarly, in some embodiments, the functionality of some of the illustrated components may not be provided and/or other additional functionality may be available.

Those skilled in the art will also appreciate that, while various items are illustrated as being stored in memory or on storage while being used, these items or portions of them may be transferred between memory and other storage devices for purposes of memory management and data integrity. Alternatively, in other embodiments some or all of the software

components may execute in memory on another device and communicate with the illustrated computer system via inter-computer communication. Some or all of the system components or data structures may also be stored (e.g., as instructions or structured data) on a computer-accessible medium or a portable article to be read from an appropriate drive, various examples of which are described above. In some embodiments, instructions stored on a computer-accessible medium separate from the computer system may be transmitted via transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as a network and/or a wireless link. Various embodiments may further include receiving, sending or storing instructions and/or data implemented in accordance with the foregoing description upon a computer-accessible medium. Accordingly, the present invention may be practiced with other computer system configurations.

Conclusion

Various embodiments may further include receiving, sending or storing instructions and/or data implemented in accordance with the foregoing description upon a computer-accessible medium. Generally, a computer-accessible medium may include storage media or memory media such as magnetic or optical media such as disks or DVD/CD-ROM, volatile or non-volatile media such as RAM, ROM, flash drives, as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as network and/or a wireless link.

The various methods described herein represent example embodiments of methods. These methods may be implemented in software, hardware, or through a combination of hardware and software. The order of the method steps may be changed, and various elements may be added, reordered, combined, omitted, or modified.

Various modifications and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. It is intended that the invention embrace all such modifications and changes and, accordingly, the above description to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A method, comprising:
 - performing, by one or more computing devices:
 - generating a representation of a map region divided into a plurality of sections;
 - receiving height information corresponding to the map region,
 - wherein the height information includes height values corresponding to one or more locations in the map region;
 - extruding each respective section of the plurality of sections according to a corresponding respective height value of the height values,
 - wherein said extruding is based on a correspondence of the respective section to a respective location in the map region, and
 - wherein the respective height value corresponds to the respective location;
 - generating a three-dimensional representation of the map region based on said extruding each respective section of the plurality of sections;
 - determining, based on the three-dimensional representation of the map region, that a label in a map view of the map region is at least partially occluded by one or more of the extruded sections below a threshold amount of partial occlusion; and

45

rendering at least a portion of the label in the map view in response to said determining that the label in the map view of the map region is partially occluded by the one or more extruded sections below the threshold amount of partial occlusion, wherein a label in the map view of the map region occluded above the threshold amount of partial occlusion is not rendered in the map view.

2. The method of claim 1, further comprising: updating the map view in response to user input indicating a change in a virtual camera viewpoint.

3. The method of claim 1, wherein subsequent to said generating the three-dimensional model, the program instructions are further executable by the at least one processor to: receive an indication of a new map region; request updated height information, wherein the updated height information overlaps with the height information; and generate another three-dimensional model based on the updated height information.

4. The method of claim 1, wherein said determining that a label in the map view of the map region is occluded below the threshold amount is further based on a virtual camera viewpoint.

5. The method of claim 1, wherein the three-dimensional model is based on correlations between two-dimensional mapping information and three-dimensional mapping information.

6. A system, comprising: at least one processor; and a memory comprising program instructions, wherein the program instructions are executable by the at least one processor to: receive mapping information corresponding to a map region, wherein the mapping information includes height values for the map region; determine, based on the height values, a three-dimensional model of the map region; determine, based on the three-dimensional model of the map region and based on coordinates for a label in the map view of the map region, a measure of partial occlusion, by one or more three-dimensional objects in the map view of the region, for the label in the map view, wherein the label in the map view is at least partially occluded by the one or more three-dimensional objects in the map view of the map region, and wherein the coordinates for the label in the map view correspond to coordinates for the label in the three-dimensional model; and render the label in the map view in response to the measure of partial occlusion being below a threshold amount of partial occlusion to be included in the map view and not render the label in response to the measure of partial occlusion exceeding the threshold amount of partial occlusion to be included in the map view.

7. The system of claim 6, wherein to determine the three-dimensional model of the map region the program instructions are further executable by the at least one processor to: extrude a plurality of sections of a grid representation of the map region, comprising: correlating one or more points of the grid representation of the map region to the height values;

46

determining a respective height value for each section of the plurality of sections based on a correlation of a respective section to a respective point in the mapping information; and extruding each section of the plurality of sections into three-dimensional space according to the respective height value in the height values corresponding to the respective point in the mapping information.

8. The system of claim 7, wherein the program instructions are further executable by the at least one processor to: adjust a granularity of the three-dimensional model based on a configuration setting indicating a level of resolution for the map, wherein an adjustment of the granularity of the three-dimensional model comprises increasing or decreasing the number of the plurality of sections.

9. The system of claim 8, wherein the program instructions are further executable by the at least one processor to: adjust one or more bases on which the measure of occlusion is determined according to the granularity of the three-dimensional model.

10. The system of claim 6, wherein the program instructions are further executable by the at least one processor to: adjust one or more bases on which the measure of occlusion is determined according to crowdsourcing information.

11. A computer-readable storage medium storing program instructions, wherein the program instructions are computer-executable to implement: receiving a set of mapping information for a map region; receiving an other set of mapping information for the map region; correlating the set of mapping information with the other set of mapping information; generating a three-dimensional model based on said correlating the set of mapping information with the other set of mapping information; calculating, based on a virtual camera viewpoint, one or more measures of partial occlusion of a label in a map view of the map region by one or more three-dimensional objects in the three-dimensional model of the map region, wherein said calculating is based on a correlation between the three-dimensional model and the map view of the map region, wherein the label in the map view of the map region is at least partially occluded by the one or more three-dimensional objects in the three-dimensional model of the map region; and rendering the label in the map view of the map region in response to calculating the one or more measures of partial occlusion of the label in the map view of the map region are below a threshold amount of partial occlusion and not rendering the label in the map view of the map region in response to calculating the one or more measures of partial occlusion of the label in the map view of the map region exceed the threshold amount of partial occlusion.

12. The computer-readable storage medium of claim 11, wherein said rendering displays less than the entire label in the map view of the map region.

13. The computer-readable storage medium of claim 11, wherein the program instructions are computer-executable to further implement: determining whether one or more areas of the label that are not occluded provide useful information, wherein said determining is based on a percentage of text in the label that is visible; and

47

wherein said calculating is further based on said determining whether the one or more areas of the label that are not occluded provide useful information.

14. The computer-readable storage medium of claim 11, wherein the virtual camera viewpoint is based on a user selection of a ground-level view or a user selection of a bird's eye.

15. The computer-readable storage medium of claim 11, wherein a measure of occlusion of the one or more measures of occlusion is based on a level of granularity of the three-dimensional model, and wherein the level of granularity of the three-dimensional model is adjustable.

16. A method, comprising:

performing, by one or more computing devices:

receiving mapping information for a map region;

approximating, based on the mapping information for the map region, a three-dimensional model of the map region;

determining a measure of partial occlusion for a label in a map view of the map region, wherein said determining is based on the three-dimensional model of the map region and based on a correlation between the three-dimensional model of the map region and the map view of the map region, wherein the label in the map view of the map region is at least partially occluded by one or more three-dimensional objects in the three-dimensional model of the map region; and rendering, based on the measure of partial occlusion for the label in the map view of the map region being less than a threshold value of partial occlusion, the label in the map view;

wherein a label in the map view of the map region partially occluded above the threshold value of partial occlusion is not rendered in the map view.

17. The method of claim 16, further comprising:

determining a level of granularity for a three-dimensional model; and

calculating the three-dimensional model according to the level of granularity,

wherein the three-dimensional model is based on the mapping information, and

wherein the three-dimensional model is based on two-dimensional mapping information and three-dimensional mapping information.

18. The method of claim 17, wherein said determining the level of granularity is based on a configuration setting indicating a level of accuracy for the occlusion testing.

19. The method of claim 18, further comprising:

adjusting one or more bases on which the measure of occlusion is based according to a configuration setting indicating a level of tolerance for partially occluded labels.

48

20. The method of claim 16, wherein said determining the measure of occlusion is based in part on data received from a mobile device within proximity of wireless communication.

21. A computer-readable storage medium storing program instructions, wherein the program instructions are computer-executable to implement:

receive mapping information for a map region;

generate a three-dimensional model based on the mapping information for the map region;

identify a label to be rendered in a map view of the map region;

determining, for the label to be rendered in the map view of the map region, label coordinates within the three-dimensional model;

determining that one or more points in the label coordinates within the three-dimensional model are unobstructed by one or more three-dimensional objects within the three-dimensional model from a virtual camera viewpoint, wherein the label is at least partially obstructed from the virtual camera viewpoint by one or more three-dimensional objects within the three-dimensional model; and

rendering, in response to the one or more points in the label coordinates within the three-dimensional model being unobstructed from the virtual camera viewpoint satisfying a threshold value for partial obstruction, at least a portion of the label in the map view, otherwise not rendering the label.

22. The computer-readable storage medium of claim 21, wherein the virtual camera viewpoint is based on coordinates, within the three-dimensional model, of a device storing and executing the program instructions.

23. The computer-readable storage medium of claim 21, wherein the program instructions are further computer-executable to implement:

generating the three-dimensional model from the two-dimensional mapping information,

wherein a respective footprint for each object in the two-dimensional mapping information is extruded to a height corresponding a default height setting for a type of a respective object.

24. The computer-readable storage medium of claim 21, wherein the program instructions are computer-executable to further implement generating the map view based on the three-dimensional model.

25. The computer-readable storage medium of claim 24 wherein the program instructions are computer-executable to further implement generating an updated map view based on input indicated a change in the virtual camera viewpoint, wherein the updated map view is based on the already generated three-dimensional model.

* * * * *